

Exhibit E

BIG-IP® Policy Enforcement Manager™: Implementations

Version 12.1



Table of Contents

Overview.....	9
What is Policy Enforcement?.....	9
About enforcement policies.....	10
About enforcement policy rules.....	10
About subscriber provisioning through PCRF.....	11
Best practices for creating enforcement policies.....	11
About sizing considerations.....	12
 Setting Up Application Visibility.....	 15
Overview: Setting up application visibility.....	15
What is application visibility?.....	15
Determining and adjusting traffic classifications.....	15
Creating a category.....	16
Creating classification presets.....	16
Creating a data plane virtual group	17
Examining application visibility statistics.....	18
 Configuring Intelligent Traffic Steering.....	 19
Overview: Configuring intelligent traffic steering.....	19
What is traffic steering?.....	19
Creating a pool.....	20
Creating forwarding endpoints.....	20
Creating an enforcement policy.....	22
Creating custom action policies.....	22
Adding rules to an enforcement policy.....	23
Creating a rule using classification criteria.....	24
Creating a rule using URL categorization.....	25
Modifying iRule event for URL categories.....	26
Creating a rule using flow conditions.....	26
Creating a rule for forwarding traffic.....	27
Creating a rule for QoS.....	28
Creating a data plane virtual group	30
 Inserting Content in to HTML Pages.....	 33
Overview: Inserting content into HTML pages.....	33
Configuring content insertion for policy action.....	33
 Configuring Quota Management using Rating Groups.....	 35

Overview: Configuring quota management	35
About Gy support and rating groups.....	35
About Online Charging System.....	36
What defines a rating group?.....	36
Configuring quota management for global settings.....	37
Creating rating groups.....	37
Adding rating group in enforcement policy.....	38
Creating a listener for quota management.....	39
About Logging Policy Enforcement Events to IPFIX Collectors.....	41
Overview: Configuring IPFIX logging for PEM.....	41
Assembling a pool of IPFIX collectors.....	41
Creating an IPFIX log destination.....	42
Creating a publisher	43
Implementation result.....	43
Reporting Usage Data to an External Analytics Server.....	45
Overview: Reporting usage data to an external analytics server.....	45
Creating a publisher	45
Creating a rule for high-speed logging for session reporting.....	46
Creating a rule for high-speed logging for flow reporting.....	47
Creating a high-speed logging rule for transactional reporting.....	48
Creating a high-speed logging rule for device detection and tethering.....	49
Session-based reporting format.....	49
Flow-based reporting format.....	51
Transaction-based reporting format.....	52
Device Type and OS-based reporting format.....	53
Tethering-based reporting format.....	54
Using Fast L4 for Performance Optimization.....	57
Overview: Using Fast L4 for performance enhancement.....	57
Creating a Fast L4 profile in PEM.....	57
Disabling a Fast L4 profile in PEM.....	58
Reporting Quality of Experience and Video Usage.....	59
Overview: Reporting quality of experience and video usage.....	59
Creating a publisher	59
Adding Quality of Experience profile to the virtual server.....	60
Configuring QoE Reporting.....	60
QoE-based reporting format.....	61
Performing Radius Authentication and Accounting.....	63
Overview: Performing RADIUS authentication and accounting.....	63

Creating a RADIUS AAA profile for policy enforcement.....	63
Creating a listener for RADIUS AAA Virtual.....	64
Creating policy rule for RADIUS accounting reports.....	64
Configuring Subscriber Discovery based on DHCP.....	67
Overview: Configuring subscriber discovery based on DHCP	67
Creating a listener for DHCPv4 discovery virtual.....	68
Creating a DHCPv4 profile for policy enforcement.....	69
Creating a listener for DHCPv6 discovery virtual.....	70
Creating a DHCPv6 profile for policy enforcement.....	71
Creating a listener for RADIUS subscriber discovery	73
Usage Monitoring Over a Gx Interface.....	75
Overview: Usage monitoring over a Gx and Sd interface	75
Creating a listener for subscriber discovery and policy provisioning.....	75
Creating a rule for usage monitoring.....	76
Configuring Global Application Policies with Bandwidth Control.....	79
Overview: Global Application Policies with Bandwidth Control.....	79
Creating VLANs.....	79
Creating a static bandwidth control policy.....	80
Creating an enforcement policy.....	81
Creating a rule for bandwidth control.....	81
Creating a listener: example	82
Managaing Traffic with Bandwidth Controllers.....	85
Overview: Bandwidth control management.....	85
Bandwidth controllers vs. rate shaping.....	85
About static bandwidth control policies.....	85
Task summary for creating a static bandwidth control policy.....	86
About dynamic bandwidth control policies.....	87
About the Maximum User Rate PPS setting.....	87
Task summary for creating a dynamic bandwidth control policy.....	87
Enforcing Bandwidth Control Provisioned by PCRF.....	91
Overview: Enforcing bandwidth control provisioned on PCRF.....	91
Creating a dynamic bandwidth control policy for PCRF.....	91
Creating a listener for subscriber discovery and policy provisioning.....	92
Implementation result.....	93
Configuring Tiered Services with Bandwidth Control.....	95
Overview: Configuring tiered services with bandwidth control.....	95
Creating dynamic bandwidth control policies for tiered services.....	95

Creating enforcement policies for three tiers.....	96
Creating the rules for tiered bandwidth control.....	97
Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF.....	98
Implementation result.....	99
Configuring Service Chains.....	101
Overview: Configuring service chains.....	101
About services profiles	102
About service chain processing	102
Creating a ICAP profile for policy enforcement.....	102
Creating a Request Adapt profile.....	103
Creating a Response Adapt profile.....	103
Creating an internal virtual server for ICAP server.....	104
Creating a pool.....	104
Creating endpoints for service chains.....	105
Creating dynamic service chains.....	105
Creating an enforcement policy.....	106
Configuring steering action policy.....	107
Adding rules to an enforcement policy.....	108
Creating a rule for forwarding traffic.....	109
Creating a data plane virtual group	110
Provisioning Dynamic Subscribers.....	113
Overview: Provisioning dynamic subscribers	113
Provisioning dynamic subscribers.....	113
Provisioning Static Subscribers.....	115
Overview: Provisioning static subscribers	115
Provisioning multiple subscribers.....	115
Subscriber CSV file format.....	116
Provisioning a file of static subscribers.....	117
About Subscriber and Policy Provisioning.....	119
Overview: Subscriber and policy provisioning	119
Configuring RADIUS in PEM profile.....	119
Configuring RADIUS AVP in PEM profile.....	120
Configuring RADIUS message.....	121
Configuring Gx in PEM profile.....	122
Configuring Gx message.....	122
Configuring Subscriber Attributes in PEM profile.....	123
Configuring Diameter AVP in PEM profile.....	124

Updating Signatures for Application Recognition.....	125
Overview: Updating classification signatures.....	125
Scheduling automatic signature updates.....	125
Creating Custom Classifications.....	127
Overview: Creating custom classifications.....	127
Determining and adjusting traffic classifications.....	127
Creating a category.....	127
Creating classification presets.....	128
Creating a custom URL database.....	128
Using iRules with classification categories and applications.....	129
Modifying iRule event for URL categories.....	130
Classification iRule commands.....	131
Configuring PEM with Local Traffic Policies.....	133
Overview: Creating local traffic policy rules for PEM.....	133
About strategies for local traffic policy matching.....	133
About creating custom local traffic policy rules for CE profile.....	134
Creating custom local traffic policy for PEM.....	135
Creating custom local traffic policy rules for PEM.....	136
Creating a virtual server for SSL traffic policy enforcement.....	136
Associating a published local traffic policy with a virtual server.....	137
Creating custom local traffic policy for PEM.....	137
Configuring Policy and RADIUS Updates.....	139
Overview: Configuring policy and RADIUS updates	139
Configuring PEM options.....	139
Creating Custom TAC DB.....	141
Overview: Custom TAC DB.....	141
Creating a custom TAC DB.....	141
Enforcing Policy and Classification on IP Protocols.....	143
About enforcing policy and classification on IP protocols.....	143
Creating Any IP profiles for PEM.....	143
Updating Any IP profile	144
IPOther filter for current PEM actions.....	144
Detecting Tethering, Device Operation System and Type.....	145
Overview: Detecting Tethering, Device Operating System, and Type.....	145
Configuring device type, OS, and tethering.....	145

Configuring PEM policy action with tethering.....	146
Creating a high-speed logging rule for device detection and tethering.....	147
Implementation result.....	148
Troubleshooting.....	149
PEM troubleshooting	149
Subscriber and policies active sessions.....	151
Active sessions statistics.....	151
Configuring subscriber activity log.....	152
IPFIX Templates for PEM Events.....	153
Overview: IPFIX templates for PEM events.....	153
About IPFIX Information Elements for PEM events.....	153
IANA-defined IPFIX Information Elements.....	153
IPFIX enterprise Information Elements.....	153
About individual IPFIX templates for each event.....	155
Session logs.....	156
Flow logs.....	156
Transaction logs.....	157
Legal Notices.....	159
Legal notices.....	159

Overview

What is Policy Enforcement?

BIG-IP® Policy Enforcement Manager™ (PEM) facilitates mobile service providers control subscriber traffic. The system can analyze application traffic and subscriber behavior, and then enforce traffic policing rules that you define. For example, you could have the system drop all web traffic coming from certain IP addresses. You can perform QoS actions on traffic you want to be treated as high priority using DSCP marking, link QoS, or bandwidth control. You could redirect HTTP traffic destined for a particular IP address, and send it to a specific URL. Or, you could send all video traffic from certain subscribers to servers for optimization.

The BIG-IP system is inserted between the subscribers and the network they are trying to access. The system intercepts the traffic that subscribers are sending. The goal of the Policy Enforcement Manager is to apply an enforcement policy to a subscriber. To determine what kind of policy to apply to the subscriber, PEM™ needs to obtain the subscriber identity.

The system can obtain subscriber identity by looking at RADIUS traffic (if present), or by analyzing subscriber data traffic. RADIUS provides much more information about the subscriber. Although analyzing subscriber traffic is more limited, it does provide the subscriber IP address. The system must have the subscriber IP address in order for PEM to do policy enforcement.

Here is a typical illustration of how policy enforcement works. Traffic from a mobile service provider goes to the BIG-IP system on its way to a network. In order to regulate subscribers, PEM needs to determine the policy to apply. For that reason, PEM collects subscriber identity by intercepting RADIUS traffic when subscriber logs in to the network and examines (snooping) the RADIUS Authentication and Accounting packets for details about the subscriber. PEM communicates with an external policy server, in this case, PCRF, for dynamic subscriber provisioning. Using the RADIUS information (or the IP address if no RADIUS is present) obtained from the subscriber identity, PEM queries the PCRF for the policy configuration and provisions subscribers dynamically.

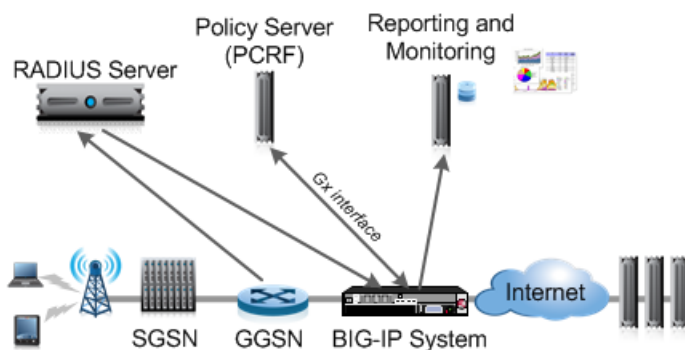


Figure 1: Diagram of policy enforcement overview

Alternatively, you can provision subscribers manually. These subscribers are called *static subscribers*. You use PEM to add static subscribers one at a time, or to import a list of subscribers. Provisioning static subscribers might require the ability to snoop RADIUS traffic but does not require a PCRF connection, as the policy assigned for static subscriber is pre-configured.

When adding static subscribers on the BIG-IP system, you provide the subscriber ID, subscriber ID type, and one or more policies to apply. You can also specify the IP address, but if it is dynamically assigned, you cannot include it. In this case, you need interception of RADIUS traffic in order to map the subscriber to the IP address. When the subscriber enters the network, the IP address from RADIUS is combined with the information already on PEM. If the static subscriber includes the IP address, no RADIUS interception is required.

About enforcement policies

An *enforcement policy* is a set of rules that determines what to do with specified types of traffic. You can configure policies on the BIG-IP® system using Policy Enforcement Manager™ (PEM), or receive policy definition from a PCRF.

For a policy to take effect, it needs to be assigned, or provisioned, to a subscriber session (a subscriber and multiple IP address mapping). A subscriber session is the period of time from when a subscriber logs into the network (authenticated) and logs out, or when the session is terminated by other means. The session is identified by subscriber IP address.

PEM™ supports the following methods for provisioning subscriber policies:

- Subscriber policy provisioning using PCRF
- Subscriber policy provisioning using a subscriber database of static subscribers (up to 100K subscribers)
- Subscriber policy provisioning for unknown subscribers (subscribers that do not currently have policies assigned to them either dynamically or statically)
- Global policy provisioning
- Custom policy provisioning using iRules®

You can use more than one of the subscriber policy provisioning methods. For example, PEM provisions an unknown subscriber policy for a subscriber session, while awaiting a response from PCRF. Or, a global policy might be applied concurrently to other subscriber policies.

As with other BIG-IP modules, like Local Traffic Manager™, you enable PEM functionality by attaching the corresponding profile to one or more virtual servers. To simplify configuration, PEM provides a listener entity that creates the required virtual servers, enables classification, and attaches the policy enforcement profile. When you create a listener, you also define which policies to apply globally or to unknown subscribers.

Advanced users can directly create virtual servers, then configure and attach the Policy Enforcement profile. We recommend that you begin configuring PEM by using listeners instead of using the advanced method. You can get familiar with PEM configuration by examining the virtual servers, settings, and profiles that the listener creates.

About enforcement policy rules

An enforcement policy is made up of a set of rules. In the policy, rules define what to do when the system receives a particular type of traffic. There are many ways you can set up a rule so that you can handle the traffic exactly as you need to. Each rule includes a condition and an action.

A rule defines conditions that the traffic must meet (or not meet) for the rule to apply. The conditions fall into the following criteria:

- *Classification criteria*, such as applications or categories of applications that the system detects. For example, a rule can apply to all webmail traffic or to a specific webmail application.

- *Flow information*, such as traffic associated with specific source and destination IP addresses or ports, or incoming DSCP marking. For example, a rule can apply to all traffic directed to a specific destination port.
- *URL information*, such as URL categories that the system detects. For example, the rule may categorize adult traffic and prevent access to it.
- *Custom criteria*, which are other conditions that you develop using iRules®

If the traffic meets the criteria in the rule, the rule specifies actions to take, such as:

- Dropping traffic
- Forwarding traffic to a specific endpoint or series of endpoints for value-added services
- Redirecting HTTP traffic to a URL
- Generating reporting data for further processing by external analytic systems
- Usage monitoring about the traffic to the PCRF so it can track mobile usage.
- Setting DSCP bits in the IP header of the traffic by marking all or marking upon the traffic exceeding a threshold
- Setting Layer 2 Quality of Service (QoS) levels for the traffic
- Enforcing rate control using a bandwidth control policy

Because rules provide so much flexibility, you need to plan what you want to do, and consider your options before you add the rules. One option is to simply classify traffic and review reports of the types of traffic your system is receiving to get more information on which to base the rules. This could be the first step when developing enforcement policies using PEM.

About subscriber provisioning through PCRF

When you are provisioning subscriber policies through PCRF, the policies are communicated using Gx interface in the form of Policy and Charging Control (PCC) rules. A PCC rule can contain the complete rule definition, or it might refer to a predefined or dynamic policy rule, as defined by the Gx protocol specification, Release 9.4. See the 3GPP TS 29.212 specification for details. When the complete rule definition is sent, it is a dynamic PCC rule; when the rule is referenced by name, it is called a predefined rule.

A *predefined PCC rule* on PCRF maps to an enforcement policy in PEM™. For example, a predefined PCC rule with the name `premium-video` on the PCRF applies to video traffic for premium subscribers. In PEM, you can create a policy also called `premium-video` with policy rules that define the enforcement action. The classification criteria for the traffic is video, and the action could be to enforce QoS for the video traffic (for example, specifying a higher bitrate).

A *dynamic PCC rule* is dynamically provisioned by the PCRF over the Gx interface. In this case, the PCC rule contains the rule definition. Therefore, in this case, you do not need to create policies on the BIG-IP® system, since the policy is totally defined on the PCRF.

Best practices for creating enforcement policies

Follow these general recommendations when creating enforcement policies:

- When creating enforcement policies you plan to apply globally or to unknown subscribers, include the word `global` or `unknown` in the policy name to distinguish these from other types of subscriber policies.
- Be cautious when developing enforcement policies to be applied globally. The policies affect all the subscribers and are applied to subscriber policies in parallel.

- When you remove or disable an enforcement policy, first be sure that it is not currently assigned to any subscribers. At least one policy must be assigned to a subscriber at all times.
- Assign the subscriber IP address when creating static subscribers that include a global or unknown subscriber policy, to ensure that the subscriber gets the entitled service faster and does not have to wait for processing of RADIUS traffic.
- If you want to use different types of steering, create separate policies and rules. For example, consider creating a policy that steers traffic from a source VLAN to an endpoint, and another policy to steer VLAN traffic to a service chain.
- Create an empty pool and use it in a forwarding endpoint if you want to route traffic or resolve policy priority conflicts between routing and steering.

These are best practices when writing policy rules:

- Be careful when you mix both L4 and L7 classification criteria in one rule; in some cases, L4 criteria takes precedence. Keep it simple: one rule, one type of criteria.
- Specify different precedence values for the rules that might conflict, to make clear in what order the rules will be evaluated.
- Do not mix different types of policy actions in the same rule; create separate rules for forwarding, reporting, Quality of Service (QoS) actions and finally, for which policy action is implemented.
- A policy (or a rule) should not direct traffic to both a forwarding endpoint and to a service chain. If both are specified, the service chain always takes precedence and is performed first, then traffic is forwarded to the endpoint.
- Dedicate certain bandwidth controllers for use only in PEM™ QoS actions, and do not use them outside PEM.
- One dynamic bandwidth controller can be applied per direction per subscriber and up to eight static bandwidth controllers can be applied, in PEM™.

There are best practices to consider when setting up reporting in enforcement policies:

- Choosing more frequent intervals for generating periodic reporting records (particularly session-based) can greatly increase the amount of reporting data, and could potentially overload the analytics system.
- Flow-based records are generated several times during the flow life and can significantly impact the amount of reporting data sent.

Here are best practices to consider when setting up iRule action:

- If multiple PEM iRules® match a flow, all the iRules are processed. The priority order is as follows:
 - PEM policy priority, and it takes in to consideration if the policy is a global high precedence policy, subscriber policy or low precedence policy.
 - PEM iRule event priority and the default event priority is 500. The event priority can be changed by specifying the priority within the iRule event.
 - Rule precedence.

About sizing considerations

Currently the maximum number of applications or category IDs that PEM™ can store, or report usage statistics for, is limited to 15 per subscriber. This in turn influences the rules, since the traffic statistics for each application or category ID is part of the rule's classification criteria.

When this limitation is exceeded for a given subscriber, an error message is logged into TMM log file. In addition, if the affected rule is installed by PCRF (over Gx connection), a session provisioning failure report is sent back to PCRF. The application or category IDs limitation should be taken into account when designing the subscriber and global policies for a particular PEM deployment.

Note: *If you require granular reporting for large amount of traffic, your performance might be impacted.*

Real performance depends on various factors such as:

- Turning on flow reporting (both performance and memory impacted)
- Frequency of sending session based reporting
- Complexity of classification
- The number of concurrent flows per subscriber
- Over subscription

Setting Up Application Visibility

Overview: Setting up application visibility

This implementation describes how to set up the Policy Enforcement Manager™ (PEM) to analyze application traffic on the network, and provide statistics for application visibility. For example, you can view statistics to see what applications are being used. By monitoring your traffic, you can later create enforcement policies that are tailored for your needs.

Task summary

Determining and adjusting traffic classifications

Creating a category

Creating classification presets

Creating a data plane virtual group

Examining application visibility statistics

What is application visibility?

Policy Enforcement Manager™ (PEM™) gives the BIG-IP® system the ability to classify both encrypted and unencrypted traffic into categories for application visibility. You can display statistics about the network traffic in graphical charts, and view classification information by application, category, protocol, virtual server, country, type of device, and so on. In-depth information and application awareness provides visibility into your network infrastructure so you can identify and monitor different types of traffic and resolve performance issues.

Application visibility is particularly useful for service providers. If your organization is using RADIUS protocol for authentication, authorization, and accounting, PEM can intercept accounting messages to retrieve additional information, for example, about mobile devices, subscribers, towers, service plans, and manufacturers.

Charts shown on the **Statistics > Classification** screens display the application visibility data. The classification overview is customizable so you can display the charts or tables that you want. The overview shows top statistics for the categories of which you are most interested.

Determining and adjusting traffic classifications

The BIG-IP® system classifies many categories of traffic and specific applications within those categories. You can determine which categories and applications of traffic the system can classify, and find out information about them such as their application or category ID.

1. On the Main tab, click **Traffic Intelligence > Classification > Applications**.
The Applications screen displays a list of the supported classification categories.
2. To view the applications in each category, click the + icon next to the category.
3. To view or edit the properties of the application or category, click the name to open its properties screen.

Tip: Here you can view the application or category ID number.

4. Click **Update** to save any changes.

Creating a category

On the BIG-IP® system, you can create customized categories for classifying traffic if the predefined categories are not sufficient for your needs. For example, if you plan to create new application types unique to your organization, you can create a category to group them together.

1. On the Main tab, click **Traffic Intelligence > Classification > Applications**.
The Applications screen displays a list of the supported classification categories.
2. Click **Create**.
The New Application screen opens.
3. From the **Type** list, select **Category**.
4. In the **Name** field, type a name for the classification category.
5. In the **Description** field, type optional descriptive text for the classification presets.
6. In the **Category ID** field, type an identifier for this category, a unique number.
7. For the **Application List** setting, move applications that you want to associate with this category from the **Unknown** list to the **Selected** list.
If the applications are not listed yet, you can associate the applications with the category when you create them.
8. Click **Finished**.

You have created custom applications to handle traffic.

Creating classification presets

On the BIG-IP® system, you can create classification preset settings for a classification policy that you have previously created.

1. On the Main tab, click **Traffic Intelligence > Classification > Presets**.
The Presets screen displays a list of the supported classification categories.
2. Click **Create**.
The New Presets screen opens.
3. In the **Name** field, type a name for the application.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. For the **Policy** setting, move the classification policies from **Available** list to the **Selected** list, to create a new preset.
6. In the **Allow Reclassification** list, **Enabled** is the default selection.
7. In the **Flow Bundling** list, **Enabled** is the default selection.
8. In the **Cache Results** list, **Enabled** is the default selection.
9. Click **Finished**.

Overview: Creating custom classifications

Creating a category

Creating a custom URL database

Creating a data plane virtual group

If you want to steer specific traffic (or otherwise regulate certain types of traffic) you must first develop appropriate enforcement policies. If using a Gx interface to a PCRF, you need to create a new virtual group in listeners that connect to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener performs preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Data Plane**.
The Data Plane screen opens.
3. Click **Add Group**.
The New Virtual Group screen opens.
4. In the **Name** field, type a unique name for the listener.
5. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

6. For the **Service Port** setting, type or select the service port for the virtual server.
7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

10. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP® system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Examining application visibility statistics

Before you can look at the application visibility statistics, you must have Adobe® Flash® Player installed on the computer where you plan to view them.

You can review charts that provide application visibility for traffic on your network.

1. On the Main tab, click **Policy Enforcement > Analytics > Overview**.
The Overview screen opens where you can view a summary of the top classification statistics.
2. Review the statistics provided. To quickly change the format of the information, click the icon to the left of the time period.
You can display information in a table, line chart, pie chart, or bar chart.
3. Click the time period (**Last Hour**, **Last Day**, **Last Week**, **Last Month**, or **Last Year**), to change the interval used for displaying content.
4. To permanently change the format or content of any of the charts, click the cog on the chart, select **Settings**, and adjust the fields in the form.
5. To display additional charts or tables, click the **Add Widget** link and complete the form.
The chart you create becomes a permanent part of the Classification Overview screen.
6. On the Main tab, click **Policy EnforcementAnalyticsStatistics**.
The Statistics screen opens and the charts display detailed classification statistics by application.
7. Adjust the statistics content in any of the following ways:
 - Use the **View By** setting or Advanced Filters to change the type of classification data shown.
 - Use the **Time Period** setting to change the interval for which statistics are shown.
 - Use the **Expand Advanced Filters** setting to fine-tune even further which types of reports to display.
8. Get detailed information in any of the following ways:
 - Point on the charts to display the details.
 - Review the Details table to see the statistics.
 - In the Details table, click the name of one of the items (application, category, protocol, and so on) to see classification details about that specific item.
 - Use the **Display method** setting to show statistics in different formats.

The easiest way to learn what classification information is available is to look at the charts and view the content and details in different ways. As you drill down into the statistics, you can locate more details and view information for a specific item.
9. To generate and export a PDF or CSV file of a report to save or email, click **Export**, select the settings, and fill in the appropriate fields.

***Note:** You must have an SMTP email server configured to use the email option. On the Main tab, click **System > Configuration > Device > SMTP**.*

You can use the classification statistics to determine, for example, the types of applications and the specific applications that clients are using. By drilling down into that information, you can find out specifically which applications are being used by a particular IP address.

Configuring Intelligent Traffic Steering

Overview: Configuring intelligent traffic steering

You can use the Policy Enforcement Manager™ to set up the BIG-IP® system to classify and intelligently steer traffic on the network. The system automatically sets up virtual servers for TCP and UDP traffic so that the BIG-IP system can classify the traffic and direct it to one or more steering endpoints based on traffic characteristics.

Note: Common Address Redundancy Protocol (CARP) persistence is supported with PEM forwarding endpoints, for use with service chaining action, when forwarding traffic to a pool.

Task Summary

- Creating a pool*
- Creating forwarding endpoints*
- Creating an enforcement policy*
- Creating custom action policies*
- Adding rules to an enforcement policy*
- Creating a rule using classification criteria*
- Creating a rule using URL categorization*
- Modifying iRule event for URL categories*
- Creating a rule using flow conditions*
- Creating a rule for forwarding traffic*
- Creating a rule for QoS*
- Creating a data plane virtual group*

What is traffic steering?

Policy Enforcement Manager™ provides the ability to intelligently steer traffic based on policy decision made using classification criteria, URL category, flow information, or custom criteria (iRule events). Steering, also called *traffic forwarding*, can help you police, control and optimize traffic.

You can forward a particular type of traffic to a pool of one or more servers designed to handle that type of traffic, or to a location closer to clients requesting a service. For example, you can send HTTP video traffic to a pool of video delivery optimization servers. You can have one policy option to classify each transaction which allows transaction aware steering. The ability to classify traffic for every transaction is called *transactional policy enforcement*. The classification per transaction is for HTTP traffic only.

You set up steering by creating an enforcement policy that defines the traffic that you want to send to a particular location or endpoint. Rules in the enforcement policy specify conditions that the traffic must match, and actions for what to do with that traffic. One of the actions you can take is to forward the traffic to a particular endpoint, called a *forwarding endpoint*.

You can create listeners to set up virtual servers and associate the enforcement policies with the traffic that is sent to them. The system also creates a Policy Enforcement profile that specifies the enforcement policy that the system uses, among other uses, for traffic steering.

Creating a pool

You can create a pool of servers that you can group together to receive and process traffic.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
5. Click **Finished**.
6. Repeat these steps for each pool you want to create.

The new pool appears in the Pools list.

Creating forwarding endpoints

Before you can create an endpoint, you need to create a pool that specifies where you want to direct the classified traffic.

To set up traffic steering, you need to create a forwarding endpoint, which specifies where to send the traffic. If you are configuring w-steering or service chains, you need to create multiple endpoints.

1. On the Main tab, click **Policy Enforcement > Forwarding > Endpoints**.
The Endpoints screen opens.
2. Click **Create**.
The New Endpoint screen opens.
3. In the **Name** field, type a name for the endpoint.
4. From the **Pool** list, select the pool to which you want to steer a particular type of traffic, for example, in a policy rule.
5. If you want to translate the destination address of the virtual server to that of the pool, from the **Address Translation** list, select **Enabled**. Otherwise, leave this setting disabled.
6. If you want to translate the original destination port to another port, from the **Port Translation** list, select **Enabled**. Otherwise, leave this setting disabled.
7. From the **Source Port** list, select the appropriate option for the source port of the connection.

Option	Description
Preserve	Maintains the value configured for the source port, unless the source port from a particular SNAT is already in use.
Preserve Strict	Maintains the value configured for the source port. If the port is in use, the system does not process the connection. Use this setting only when (1) the port is configured for UDP traffic; (2) the system is configured for nPath routing or running in transparent mode; or (3) a one-to-one relationship exists between

Option	Description
	virtual IP addresses and node addresses, or clustered multi-processing (CMP) is disabled.
Change	Specifies that the system changes the source port.

8. To specify a SNAT pool for address translation, from the **SNAT Pool** list, select the name of an existing SNAT pool.
The steering endpoint uses the SNAT pool to implement selective and intelligent SNATs.
9. If you have multiple pool members and want specific traffic to go to the same pool member every time, from the **Persistence** list, select the appropriate IP address type:

Option	Description
Hash Settings	Map the hash value to a specific pool member so that other traffic, with the same hash value, is directed to the same pool member.
Source Address	Map the source IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.
Destination Address	Map the destination IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.

If you do not need to maintain persistence, leave **Persistence** set to **Disabled**, the default value.

10. If you select **Hash Settings**, configure the following fields:

- To specify a algorithm for the hash persistence method, from the **Hash Persistence Algorithm** list, select the name of an algorithm. The CARP algorithm is the only options available currently.
- In the **Hash Persistence Offset** field, type the offset from start of the source string to calculate the hash value. The default value is **0**.
- In the **Hash Persistence Length** field, type the length of the source string used to calculate the hash value. The default value is **1024**.

11. From the **Hash source** list, select the appropriate method to get the hash value.

Option	Description
URI	Specify the string value to calculate hash value.
Execute Script	Specify the script for TCL script snippet. You can select the Wrap Area Text check box to wrap the definition text, and select the Extend Area check box to increase the field space of format scripts.

The results from this script are used to calculate the hash value.

***Note:** The URI option is for HTTP traffic only.*

12. If you want to apply fallback persistence method that is applied when default persistence fails, from the **Fallback Persistence** list, select the appropriate IP address type:

Option	Description
Disabled	Disables fallback persistence. The default value is Disabled .
Source Address	Map the source IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.
Destination Address	Map the destination IP address to a specific pool member so that subsequent traffic from this address is directed to the same pool member.

13. Click **Finished**.

You can direct traffic to the endpoint you created in the policy rules of an enforcement policy.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.

***Tip:** When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.*

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

***Important:** The system performance is significantly affected, depending on complexity of the classification and the type of policy action.*

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Creating custom action policies

In an enforcement policy, custom action can be defined by a Policy Enforcement Manager™ (PEM™) iRule.

1. On the Main tab, click **Policy Enforcement > Policies > iRules**.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a name for the new iRule.
4. In the **Description** field, type a description of the new iRule.
5. In the **iRule Expression** field, specify the TCL syntax that defines a custom iRule action, which can be later attached to a policy enforcement rule.

```
when PEM_POLICY { if {[PEM::policy initial]}
    { /* Commands to run during the first time the policy is evaluated.
      */
      } else
      { /* Commands to run during policy re-evaluation. */ }
    /* Commands to run during policy eval and re-eval time. */ }
```

There can be two iRule events:

- `PEM_POLICY` is triggered when a policy evaluation occurs.
- `RULE_INIT` runs the first time the iRule is loaded or has changed.

The two new PEM iRule commands are `PEM::policy initial` and `PEM::policy name`. You can select the **Wrap Text** check box to wrap the definition text, and select the **Extend Text Area** check box to increase the field space of format scripts.

6. Click **Finished.**

The Policy Enforcement Manager creates a new iRule, and displays the iRule list.

7. To attach a custom action to a specific iRule, follow the steps:

- a) Click **Policy Enforcement > Policies**.
- b) Select a policy name.
- c) Click a policy rule.
- d) From the **Custom Action** list, select a iRule created.

8. Click **Update.**

You have created a custom action in a policy, using iRules.

Note: The iRule actions are executed at the end of all the other policy actions.

Adding rules to an enforcement policy

Before you can add rules to an enforcement policy, you need to create the policy, then reopen it.

You add rules to an enforcement policy to select the traffic you want to affect, and the actions to take. A *rule* associates an action with a specific type of traffic. So you can, for example, add a rule to select all audio-video traffic and send it to a pool of servers that are optimized to handle that type of traffic.

- 1. On the Main tab, click **Policy Enforcement > Policies**.**
The Policies screen opens.
- 2. Click the name of the enforcement policy you want to add rules to.**
The properties screen for the policy opens.
- 3. In the Policy Rules area, click **Add**.**
The New Rule screen opens.
- 4. In the **Name** field, type a name for the rule.**
- 5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.**

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

- 6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.**
- 7. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.**
More modify header configuration options display.
- 8. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.**
Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

9. From the **Congestion Detection** list, select **Enable**, to congestion detection in the Radio Access Network.
 - a) In the **Threshold** field, type the lower threshold bandwidth for a session. The default value is 1000 kbs.
 - b) For **Destination** list, select the publisher name from the HSL publisher drop-down list.

The state of congestion detection is now controlled by policy application, and different subsets of subscribers can have different settings. This enables congestion-detection for specific types of applications as it pairs with specific policy rule conditions.

10. Click **Finished**.
11. Repeat steps 3-8 to create as many rules as needed to handle the traffic you are interested in.

The enforcement policy includes the rules with the conditions and actions you added.

Now you need to associate the enforcement policy with the virtual server (or servers) to which traffic is directed.

Creating a rule using classification criteria

You can use Layer 7 classification criteria to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. On the Classification tab, in the **Classification** setting, specify Layer 7 matching criteria for the rule:
 - a) From the **Match Criteria** list, select whether you want perform actions on traffic that matches (select **Match**), or does not match (select **No Match**) the criteria specified.
 - b) From the **Category** list, select the type of traffic this rule applies to, or select **Any** for all traffic.
 - c) Some categories have specific applications associated with them. If this one does, from the **Application** list select the application this rule applies to, or select **Any** for all traffic in this category.
 - d) Click **Add** to add this match criteria to the classification.
Add as many matching criteria as are relevant to this rule.

7. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

8. Click **Finished**.

You have created a rule that applies to traffic based on classification criteria.

Creating a rule using URL categorization

You have the ability to enforce policies that are configured as part of the subscriber profile, based on the URL category type. Use Layer 7 criteria to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. On the URL tab, in the **URL** setting, specify Layer 7 matching criteria for the rule :
 - a) From the **Match Criteria** list, select whether you want perform actions on traffic that matches (select **Match**), or does not match (select **No Match**) the criteria specified.
 - b) From the **URL Category** list, select the type of traffic this rule applies to.
 - c) Click **Add** to add this match criteria to the classification.
Add as many matching criteria as are relevant to this rule.

7. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

8. Click **Finished**.

You have created a rule that applies to traffic based on URL Category.

Modifying iRule event for URL categories

On the BIG-IP® system, you can modify iRules® Event settings for URL categories.

1. On the Main tab, click **Traffic Intelligence > URL Categorization > Categories**.
2. Select a URL category.
The URL Properties screen opens.
3. In the **Name** field, type a unique name for the URL category policy.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. In the **Category ID** field, type an identifier for this category, a unique number.
6. In the **iRule Event** field, select the appropriate setting.
 - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
 - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

Note: *CLASSIFICATION::DETECTED is the only event that is supported.*

You have modified an iRule event setting for an existing URL category.

Creating a rule using flow conditions

You can use flow information to define conditions that the traffic must meet (or not meet) for an enforcement policy rule to apply.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: *All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. On the Flow tab, in the **Flow** setting, specify Layer 4 conditions that the traffic must meet (or not meet) for this rule to apply.

Option	Description
Match	Select whether you want to perform actions on traffic that matches (select Match) or does not match (select No Match) the criteria specified.
DSCP Marking	To match incoming traffic based on a DSCP value, type an integer from 0 to 63.
Protocol	To specify the applicable traffic by protocol, select UDP , TCP , or leave the default value of Any .
IP Type	To specify the IP address type that this rule applies to, select IPv4 , IPv6 , or leave the default value of Any .
Source Address/Mask	To match incoming traffic based on the address or network it is coming from, type the source IP address/netmask of the network you want the rule to affect. The default value is 0.0.0.0/32.
Source Port	To match incoming traffic based on the port it is coming from, type the port number you want the rule to affect. The default value (empty) matches traffic from all ports.
Source VLAN	To match incoming traffic based on the VLAN, select a previously configured VLAN.
Destination Address/Mask	To match traffic based on the address or network it is directed to, type the source IP address/netmask of the network you want the rule to affect. The default value is 0.0.0.0/32.
Destination Port	To match incoming traffic based on the port it is directed to, type the port number you want the rule to affect. The default value (empty) matches traffic headed to all ports.

- a) Click **Add** to add this match criteria to the classification.

Tip: F5® recommends that you keep the matching criteria in a rule simple, adding more rules to specify additional conditions rather than including too many in one rule.

7. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.

Other tasks describe how to do this in detail.

If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.

8. Click **Finished**.

You have created a rule that classifies traffic.

Creating a rule for forwarding traffic

You can create a rule that forwards traffic to an endpoint. For example, you might want to direct video traffic to a server that is optimized for video viewing.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.

The properties screen for the policy opens.

3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. In the Gate area, for **Gate Status**, select **Enabled**.
Options provide several ways to forward the traffic.
8. In the Forwarding area, for **HTTP Redirect**, select **Enabled**, and type the URL.
9. From the Forwarding list, select an option where you would like to forward the traffic.

Options	Description
Route to Network	The traffic flow is forwarded to the default destination.
Forwarding to Endpoint	The flow is steered to a different destination and you can select one of the endpoints.
Forward to ICAP virtual Server	The flow is forwarded to the ICAP virtual server.

10. From the **Forwarding Fallback Action** list, select **Drop** or **Continue** to specify if the connection can remain unchanged or should be dropped if the forwarding action fails.
11. From the **ICAP Virtual Server** list, select an internal virtual server that you have created, or click **Create** to create a new internal virtual server.
12. From the **ICAP Type** list, select an ICAP adaptation type.
 - Select **Request** to send a portion of the request to the ICAP server.
 - Select **Response** to receive a portion of the response from the ICAP server.
 - Select **Request** and **Response** to have both types of adaptation.
13. From the **Service Chain** list, select **Create** to direct traffic to more than one location (such as value-added services).
14. Click **Finished**.

You have created a rule that forwards traffic.

Creating a rule for QoS

Before you can create a rule for Quality of Service (QoS), you need to create a bandwidth controller to use rate control.

You can create a rule that results in a QoS action such as DSCP marking, link QoS, or rate limiting.

Note:

In the mobile market, uplink and downlink is sometimes known as forward and reverse respectively.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. For **Gate Status**, select **Enabled**.
If you select **Disabled**, then the corresponding traffic will be dropped.
Forwarding and QoS options are displayed.
8. To set DSCP bits on the downlink traffic, for **IP Marking (DSCP)**, select **Specify**, and type a value between 0 and 63, inclusive.
The traffic that matches this rule is marked with this value.
9. To set DSCP bits on the uplink traffic, for **IP Marking (DSCP)**, select **Specify**, and type a value between 0 and 63, inclusive.
The traffic that matches this rule is marked with this value.
10. To set a Layer 2 Quality of Service (QoS) level in downlink packets, for **L2 Marking (802.1p)**, select **Specify**, and type a value between 0 and 7, inclusive.
Setting a QoS level affects the packet delivery priority.
11. To set a Layer 2 Quality of Service (QoS) level in uplink packets, for **L2 Marking (802.1p)**, select **Specify**, and type a value between 0 and 7, inclusive.
Setting a QoS level affects the packet delivery priority.
12. To apply rate control to downlink traffic, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

Note: You can assign any previously created static or dynamic bandwidth control policies. However, F5® does not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create to enforce dynamic QoS settings provisioned by the PCRF.

Depending on the bandwidth control policy, PEM™ restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, or any combination of these.

13. To apply rate control to uplink traffic and per category of application, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

Note: You can assign any previously created static or dynamic bandwidth control policies. However, we do not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create for communicating with the PCRF.

Depending on the bandwidth control policy, PEM restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, per category of applications or any combination of these.

14. Click Finished.

You have created a rule that manages QoS traffic.

Creating a data plane virtual group

If you want to steer specific traffic (or otherwise regulate certain types of traffic) you must first develop appropriate enforcement policies. If using a Gx interface to a PCRF, you need to create a new virtual group in listeners that connect to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener performs preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Data Plane**.
The Data Plane screen opens.
3. Click **Add Group**.
The New Virtual Group screen opens.
4. In the **Name** field, type a unique name for the listener.
5. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

6. For the **Service Port** setting, type or select the service port for the virtual server.
7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

10. Click Finished.

The Policy Enforcement Manager creates a listener.

When you create a listener, Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP® system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Inserting Content in to HTML Pages

Overview: Inserting content into HTML pages

You can use the Policy Enforcement Manager™ (™ PEM) for header modifications by configuring insertion of content into HTML web pages, based on a PEM policy. The content inserted can be a string or TCL script, and can occur in the beginning, middle, or end of the HTML page.

Insertion of content is useful in these scenarios:

- To insert tailor made ads to certain subscribers based on profiles.
- For sending alerts to the subscriber when the usage quota is reached.
- For periodically sending the subscriber usage statistics.

PEM inserts the content at the match tag, and, depending on the throttling level, implements:

- If the content needs to be inserted once per subscriber.
- Periodic insertions into the transactions, which matches the criteria. (You can configure the periodic timer.)
- Insertion in all the transactions matching the action.

Configuring content insertion for policy action

Before you can insert content in to the HTML pages, on the BIG-IP® system, you need to create an enforcement policy and add rules to an enforcement policy. You can add the content action as one of the actions to the existing list of actions supported by PEM. You can also control the insertion frequency through configuration per transaction.

Note: The size of the content being inserted cannot be more than 65k letters.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for

the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Insert Content** list, select **Enabled**, to modify the content insertion.

***Note:** The BIG-IP system does not allow matching tags that are inserted, by insert action, on a given flow.*

More content insertion configuration options display.

8. To insert content, select the action you want to implement.
 - From the **Position** list, select **Before** to prepend (opening and closing) the HTTP tag. Specify the tag name from the **HTTP Tag** list.
 - From the **Position** list, select **After** to append (opening and closing) the HTTP tag. Specify the tag name from the **HTTP Tag** list.
9. In the **Type** field, type the string that defines the content, or the TCL script, which can be later attached to a policy enforcement rule.

```
HTTP::header replace Server \"Nginx\"
```

There can be two iRule events:

- PEM_POLICY is triggered when a policy evaluation occurs.
- RULE_INIT runs the first time the iRule is loaded or has changed.

The two new PEM iRule commands are PEM::policy initial and PEM::policy name. You can select the **Wrap Text** check box to wrap the definition text, and select the **Extend Text Area** check box to increase the field space of format scripts.

10. From the **Frequency** list, select a frequency type.
 - Select **Always** to insert content in each transaction.
 - Select **Once** to insert content once per action.
 - Select **Once Every** to insert content once per designated time period, in seconds.
11. Click **Finished**.
12. On the Main tab, click **Policy Enforcement > Options**.
The Options screen opens.
13. For the **Content Insertion Options** setting, in the **Throttling** fields, type the time used to set the maximum wait time before Policy Enforcement Manager™ applies the insert action again on the same subscriber.

***Note:** The insert actions do not conflict with each other.*

You have inserted content in to the HTML pages.

Configuring Quota Management using Rating Groups

Overview: Configuring quota management

You can use the Policy Enforcement Manager™ to implement quota management process for prepaid subscribers per session and per application. You can provision prepaid charging per subscriber or application that communicates with the quota protocol endpoint (QPE), such as online charging system (OCS), over the 3GPP Gy interface. The Gy endpoint allows online credit control for Layer 4 to 7 service data flow-based charging. This type of policing is called quota management; this feature ensures that subscribers do not consume resources that are not authorized. The Diameter Credit-Control Application (DCCA) specifies an approach based on a series of interrogations, that use Credit-Control-Request (CCR) and Credit-Control-Answer (CCA) messages.

If threshold for a Quota bucket is almost reached and the current flow 8KB interval is not used, then CCR-U is sent after 8KB consumption (maximum). PEM forwards it and after 8KB interval, or in between if in middle, it sends CCR-U. Consider 8KB consumption intervals when Quota bucket is being checked for exhaustion.

If larger bytes arrives between 8KB interval then it is let through. After we detect Quota is being exhausted, after usage of 8KB interval, then we update the internal policy change for the session. After reevaluation interval (20 seconds) the flow is reevaluated and then termination action is applied. However, for new flows as soon as Quota expires they are dropped.

The PCRF reports the quota usage under a number of circumstances and notifies the server through the use of the Reporting-Reason AVP in the CCR. The reason for reporting credit usage can occur directly in the Multiple-Services-Credit-Control AVP or in the Used-Service-Units AVP and depends on whether it applies for all quota types or a particular quota type.

Task summary

About Online Charging System

What defines a rating group?

Configuring quota management for global settings

Creating rating groups

Adding rating group in enforcement policy

Creating a listener for quota management

About Gy support and rating groups

The Gy interface in 3GPP architecture facilitates communication between the online charging system (OCS) and the PCEF. In turn, this communication supports the advanced credit authorization and quota-specific reporting. Policy Enforcement Manager™ provides online credit control, through user configuration, for Layer 4 to 7 service data flow-based charging.

The subscriber traffic contains consumed based on allocated quota that is based on applications, category, or a group of them and is measured in terms of volume, time, and events. A rating group, which is the same as a quota bucket, can be created. A rating group is identified by a service-identifier AVP that gathers a set of services, which has the same costs and rating type. Once you create a rating group, you can assign it to multiple rules inside the policy. For all the traffic matching the rule, quota is consumed from this bucket to

make sure there is no over-subscription of resources. For example, you can have a rating group assigned to managing video traffic of 500 MB. This rating group needs to be assigned to a rule that matches the video traffic, to ensure that there is no over-subscription of subscriber video traffic.

Note: You need to assign a default rating group on your policy rule or assign a new one. The default rating group is for all traffic that does not belong to another rating group.

Multiple Services Credit Control (MSCC) is a procedure that allows quota management for multiple services within one Gy Session. It is possible to allocate quota on a per service basis or the services can be grouped into rating group to gather aggregation of quota management. When the MSCC is present in the CCR message, it represents PCRF. The MSCC requests quota for a particular service, or multiple services, or usage being reported. When the MSCC is present in the CCA message, it represent OCS that grants quota for the service or services.

About Online Charging System

Online charging uses IETF Diameter Credit Control application. It uses the Gy reference point in the 3GPP standards for messages between the Online Charging Server and the PCEF. The PCEF requests resource allocation and reports credit control information to the OCS (Online Charging System). Event Charging with Unit Reservation (ECUR) using CCR Initial and Termination can be used to generate event messages.

Online charging has two sub-functions:

- Unit determination, that refers to the calculation of the number of non-monetary units like service units, data volume, time and events that shall be assigned prior to starting service delivery. PEM uses the Session Charging with Unit Reservation (SCUR) system.
- Rating refers to determining a price, based on the non-monetary units calculated by the unit determination function. When online charging is used in the P-GW, the credit control is per rating group. P-GW allows reporting of the service usage per rating group or per combination of the rating group and service ID. Service ID sometimes is synonymous of rating group. Gy messages does not have to include both service identifier and rating group. You can use both service identifier and rating group, to be able to rate based on rating group but keep track of usage per service. For example, if you rate Facebook, Twitter and YouTube in the same way, so to create a single rating group, you accumulate usage per each service. With this setup you get less reservation spreading, while keeping track of usage per service. This reporting level can be activated per PCC rule.

Overview: Configuring quota management

Overview: Configuring quota management

What defines a rating group?

What defines a rating group?

The four Attribute Value Pair (AVP) that defines a rating group, are the amount of time, the downloaded, uploaded and total bytes.

Note: Rating group can be either the Granted Units (GSU) that is used before releasing the service or a new CCR needs to be sent, or Used Service Units (USU) which reports the total traffic for subscriber or service.

You can define rating group as follows:

- Define the rating group for thresholds and timers
- Validity time, which indicates the time that limits the validity of the granted quota

- Quota holding indicates the number of seconds for which the quota granted by the OCS, is held by the gateway when no traffic is received for that rating group.
- Quota consumption, used by the OCS, indicates to the client that the quota consumption must be stopped after a period at session termination or when no packets are received.
- Time quota threshold indicates the threshold in seconds when the granted quota is time.
- The Volume-Quota-Threshold AVP indicates the threshold in octets when the granted quota is volume.

Overview: Configuring quota management

About Online Charging System

Configuring quota management for global settings

Configuring quota management for global settings

You can set up global configuration for quota management in the BIG-IP® system. A default rating group defines the case when a classified flow does not have an quota action. This can be used in cases were PEM might have policies to do quota management for specific applications, and no policy for the remaining application.

Note: *If a rating group is configured as default, that group cannot be used by any rules.*

1. On the Main tab, click **Policy Enforcement > Options**.
The Options screen opens.
2. In the Quota Management Options area, for the **Default Rating Group** setting, select **Create** to create a new rating group for quota management.
This takes you to the **Policy Enforcement > Rating Groups > New Rating Group** screen. Click **Policy Enforcement > Options** to go back to options screen.

You created a default rating group that enables the possibility to implement quota per subscriber. Instead of defining all the policies with the corresponding rating group, you can define a single rating group and add it as a default.

Overview: Configuring quota management

What defines a rating group?

Creating rating groups

Creating rating groups

You can assign a rating group to a rule and attach it to a policy. For example, if you want to allocate quota to all the videos a subscriber uses from multiple on-demand Internet streaming media, you can specify a quota bucket that covers all the quota consumption and ensures that the consumption does not exceed the specified time or volume.

1. On the Main tab, click **Policy Enforcement > Rating Groups**.
The Rating Groups List screen opens.
2. Click **Create**.
The New Rating Group screen opens.
3. In the **Name** field, type a name for the rating group.
4. In the **Description** field, type optional descriptive text for the rating group.
5. In the **Rating Group ID** field, type an unique identifier (integer). This Rating Group ID is used by the quota managing endpoint, such as, Gy.

6. In the **Initial Quota** setting, specify **Volume** in octets, the initial quota to receive and send from the OCS, and the total quota volume.
7. In the **Initial Quota** setting, specify **Time** in seconds, the initial time for quota.
8. In the **Default Quota** setting, specify **Volume** in octets, the default quota to receive and send from the OCS, and the total quota volume.
9. In the **Threshold** field, type a default threshold level you want to use for a sending quota replenishment request.
The default value is 0, which indicates that there is no threshold.
10. In the **Usage Time** field, type the quota for how long the traffic can be used.
11. In the **Consumption Time** field, type the maximum idle time that is accounted as quota usage. This is the default value of quota for time and specifies time units for charging as well.
12. In the **Validity Time** field, type the duration for which the quota is used, if the online charging system (OCS) does not specify the validity time.
13. In the **Holding Time** field, type the holding time (in seconds), for which the quota is valid without any usage, if the time is not specified by the OCS.

***Note:** The default values for consumption time, validity time and holding time are used, when the OCS does not specify them.*

14. From the **Breach Action** list, select the appropriate action to be taken when default quota expires or OCS does not provide new quota or breach action.

Breach Action	Description
Terminate	The system stops traffic when quota is breached.
Allow	The system allows traffic to go through even when the quota is breached.
Redirect	The system redirects traffic to the forwarding endpoint, when quota is breached.

15. From the **Request on Install** list, select **Yes** if the quota has to be requested from the Gy, when the policy is installed for a subscriber. Otherwise, select **No** for quota to be requested when one of the applications associated with the rating group is detected.

Adding rating group in enforcement policy

Before you can add rules to an enforcement policy, you need to create the policy, then reopen it.

You add rules to an enforcement policy to select the traffic you want to affect, and the actions to take. A *rule* associates an action with a specific type of traffic. So you can, for example, add a rule to select all audio-video traffic and send it to a pool of servers that are optimized to handle that type of traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.

5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. From the **Rating group** list, in the **Quota** setting, select the prior configured rating group or create a new rating group. This specifies what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.
7. Click **Finished**.
8. Repeat steps 3-8 to create as many rules as needed to handle the traffic you are interested in.

The enforcement policy includes the rules with the conditions and actions you added.

Now you need to associate the enforcement policy with the virtual server (or servers) to which traffic is directed.

[Overview: Configuring quota management](#)

[Creating rating groups](#)

[Creating a listener for quota management](#)

Creating a listener for quota management

You can create listeners that specify how to handle traffic for policy enforcement. You can add support for Gy over a diameter connection by adding a diameter profile which has Gy application support to the diameter virtual server in PEM.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. In the Policy Provisioning and Online Charging Virtuals area, click **Add**.
The New Configure Diameter Endpoint Provisioning and Online Charging screen opens.
3. In the **Name Prefix** field, type a unique name for the listener.
4. In the **Description** field, type a description of the listener.
5. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
6. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
7. To connect to a PCRF, from the **Diameter Endpoint Provisioning** list, select **Gy** from the **Supported Apps** options.
8. In the **Product Name** field, type the product name which is used to communicate with the OCS.
9. In the **Origin Host** field, type the fully qualified domain name of the OCS, for example, `ocs.xnet.com`.
10. In the **Origin Realm** field, type the realm name or network in which the OCS resides, for example, `xnet.com`.
11. In the **Destination Host** field, type the destination host name of the OCS, for example, `ocsdest.net.com`.

12. In the **Destination Realm** field, type the realm name or network of the OCS, for example, `net.com`.
13. For the **Pool Member Configuration** setting, add the OCS servers that are to be members of the Gy endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
14. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the OCS over the Gy interface. The default value is 1500.
15. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the OCS. The default value is 2.
16. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
17. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system enables classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gy interface with a diameter endpoint profile.

About Logging Policy Enforcement Events to IPFIX Collectors

Overview: Configuring IPFIX logging for PEM

You can configure the BIG-IP® system to log information about Policy Enforcement Manager™ (PEM™) processes and send the log messages to remote IPFIX collectors.

The BIG-IP system supports logging of PEM events over the IPFIX protocol. IPFIX logs are raw, binary-encoded strings with their fields and field lengths defined by IPFIX templates. *IPFIX collectors* are external devices that can receive IPFIX templates and use them to interpret IPFIX logs.

The configuration process involves creating and connecting the following configuration objects:

Object	Reason
Pool of IPFIX collectors	Create a pool of IPFIX collectors to which the BIG-IP system can send IPFIX log messages.
Destination	Create a log destination to format the logs in IPFIX templates, and forward the logs to the IPFIX collectors.
Publisher	Create a log publisher to send logs to a set of specified log destinations.

Task summary

Perform these tasks to configure IPFIX logging of PEM processes on the BIG-IP® system.

Note: *Enabling IPFIX logging impacts BIG-IP system performance.*

Assembling a pool of IPFIX collectors

Creating an IPFIX log destination

Creating a publisher

Assembling a pool of IPFIX collectors

Before creating a pool of IPFIX collectors, gather the IP addresses of the collectors that you want to include in the pool. Ensure that the remote IPFIX collectors are configured to listen to and receive log messages from the BIG-IP® system.

These are the steps for creating a pool of IPFIX collectors. The BIG-IP system can send IPFIX log messages to this pool.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.

4. Using the **New Members** setting, add the IP address for each IPFIX collector that you want to include in the pool:
 - a) Type the collector's IP address in the **Address** field, or select a node address from the **Node List**.
 - b) Type a port number in the **Service Port** field.
By default, IPFIX collectors listen on UDP or TCP port 4739 and Netflow V9 devices listen on port 2055, though the port is configurable at each collector.
 - c) Click **Add**.
5. Click **Finished**.

Creating an IPFIX log destination

A log destination of the **IPFIX** type specifies that log messages are sent to a pool of IPFIX collectors. Use these steps to create a log destination for IPFIX collectors.

1. On the Main tab, click **System > Logs > Configuration > Log Destinations**.
The Log Destinations screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this destination.
4. From the **Type** list, select **IPFIX**.
5. From the **Protocol** list, select **IPFIX** or **Netflow V9**, depending on the type of collectors you have in the pool.
6. From the **Pool Name** list, select an LTM[®] pool of IPFIX collectors.
7. From the **Transport Profile** list, select **TCP**, **UDP**, or any customized profile derived from TCP or UDP.
8. The **Template Retransmit Interval** is the time between transmissions of IPFIX templates to the pool of collectors. The BIG-IP system only retransmits its templates if the **Transport Profile** is a **UDP** profile.
An *IPFIX template* defines the field types and byte lengths of the binary IPFIX log messages. The logging destination sends the template for a given log type (for example, NAT44 logs or customized logs from an iRule) before sending any of those logs, so that the IPFIX collector can read the logs of that type. The logging destination assigns a template ID to each template, and places the template ID into each log that uses that template.

The log destination periodically retransmits all of its IPFIX templates over a UDP connection. The retransmissions are helpful for UDP connections, which are lossy.
9. The **Template Delete Delay** is the time that the BIG-IP device should pause between deleting an obsolete template and re-using its template ID. This feature is helpful for systems that can create custom IPFIX templates with iRules.
10. The **Server SSL Profile** applies Secure Socket Layer (SSL) or Transport Layer Security (TLS) to TCP connections. You can only choose an SSL profile if the **Transport Profile** is a **TCP** profile. Choose an SSL profile that is appropriate for the IPFIX collectors' SSL/TLS configuration.

SSL or TLS requires extra processing and therefore slows the connection, so we only recommend this for sites where the connections to the IPFIX collectors have a potential security risk.
11. Click **Finished**.

Creating a publisher

A publisher specifies where the BIG-IP® system sends log messages for IPFIX logs.

1. On the Main tab, click **System > Logs > Configuration > Log Publishers**.
The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. Use the Log Destinations area to select an existing IPFIX destination (perhaps along with other destinations for your logs): click any destination name in the **Available** list, and click << to move it to the **Selected** list.
5. Click **Finished**.

Implementation result

Now you have an implementation in which the BIG-IP® system logs messages about PEM™ session, flow and transaction reporting and sends the log messages to a pool of IPFIX collectors.

Reporting Usage Data to an External Analytics Server

Overview: Reporting usage data to an external analytics server

In Policy Enforcement Manager™, you can create a rule within an enforcement policy that instructs the system to send usage data in high-speed logging (HSL) format to an external analytics server. The rule specifies what type of reporting data you are interested in; one of the actions it can take with the traffic is to send the information collected about it for processing to a centralized analytics server.

The system sends the information as a set of comma-separated values by means of SYSLOG transport. You can choose to use the session-based, flow-based or transactional reporting format, depending on the level of granularity you need.

For example, a rule might collect session-based information about all audio and video traffic. You can specify how often to log the data and set the destination as an HSL server or pool.

Transactional Policy Enforcement, provides the ability to report each of the HTTP transaction and sends the report to a HSL publisher. Each transaction report information is specific to that transaction only. The transactional reports are used for analytics and high level granularity for application and subscriber visibility.

Task summary

Creating a publisher

Creating a rule for high-speed logging for session reporting

Creating a rule for high-speed logging for flow reporting

Creating a high-speed logging rule for transactional reporting

Creating a high-speed logging rule for device detection and tethering

Creating a publisher

Before you create a publisher, you have to create a HSL pool that needs to be associated to a destination. Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System > Logs > Configuration > Log Publishers**.
The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click << to move the destination to the **Selected** list.

***Note:** If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

Creating a rule for high-speed logging for session reporting

Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent.

In an enforcement policy, a rule can specify that session statistics about the traffic affected by the rule are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Usage Reporting** list, select **Enabled**.
8. From the **Usage Report Granularity** list, select **Session** to log details about subscribers and application sessions.
9. In the **Usage Volume Threshold** setting, specify in octets, the threshold to send HSL reporting records. You can send reporting data from uplink traffic, to downlink traffic and the total traffic volume before logging the information.
10. In the **Usage Destination** setting, specify where to send the usage monitoring data:
 - In the **Gx** field select **Enabled** for the BIG-IP system to send usage monitoring data over a Gx interface. You can then type a string for the **Gx Monitoring Key** that is used for usage monitoring.

***Note:** When you select **Session** in the **Report Granularity** field, the **Gx** field appears.*

- From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs and select the format script of the report from the **Format Script** list.
- Select the **RADIUS Accounting** option from the destination. From the **RADIUS AAA Virtual** list, select the RADIUS AAA virtual that you have created before.

***Note:** If you are using a formatted destination, select the publisher that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

***Note:** There are no **Format Scripts** for transactional reporting.*

11. In the **Usage Interval** field, type an integer that specifies how frequently HSL reporting data is sent.
12. For the **Session Reporting Field** setting, move the fields that you want to see in the logs from the **Available** list to the **Selected** list.
13. Click **Finished**.

You have created a rule that sends data about the traffic to external high-speed logging servers. The CSV reporting format differs depending on whether the report granularity is flow-based or session-based.

Creating a rule for high-speed logging for flow reporting

Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent.

In an enforcement policy, a rule can specify that flow statistics about the traffic affected by the rule are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Usage Reporting** list, select **Enabled**.
8. From the **Report Granularity** list, select **Flow**, for more granular reporting of every TCP connection.
9. In the **Volume Threshold** setting, specify in octets, the threshold to send HSL reporting records. You can send reporting data from uplink traffic, to downlink traffic and the total traffic volume before logging the information.
10. In the **Interval** field, type an integer that specifies how frequently HSL reporting data is sent.
11. In the **Destination** setting, specify where to send the usage monitoring data:
 - From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs.
 - Select the **Format Script** list and select the format script of the report from the **Format Script** list.
 - Select the **RADIUS Accounting** option from the destination. From the **RADIUS AAA Virtual** list, select the RADIUS AAA virtual that you have created before.

***Note:** If you are using a formatted destination, select the publisher that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

12. For the **Flow Reporting Field** setting, move the fields that you want to see in the logs from the **Available** list to the **Selected** list.
13. Click **Finished**.

You have created a rule that sends data about the traffic to external high-speed logging servers. The CSV reporting format differs depending on whether the report granularity is flow-based or session-based.

Creating a high-speed logging rule for transactional reporting

Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent.

In an enforcement policy, a rule can specify that transactional statistics about traffic affected by the rule are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
3. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
4. In the Policy Rules area, click **Add**.
The New Rule screen opens.
5. In the **Name** field, type a name for the rule.
6. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

7. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
8. From the **Usage Reporting** list, select **Enabled**.
9. From the **Report Granularity** list, select **Transaction**, for more granular reporting of every HTTP transaction.
10. In the **Additional HTTP Information** setting, specify in bytes, the HTTP **Hostname**, the HTTP **User Agent**, and the HTTP **URI**.
11. In the **Destination** setting, specify where to send the usage monitoring data:
 - From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs.
 - Select the **RADIUS Accounting** option from the destination. From the **RADIUS AAA Virtual** list, select the RADIUS AAA virtual that you created earlier.

Note: If you are using a formatted destination, select the publisher that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.

Note: There are no **Format Scripts** for transactional reporting.

12. For the **Transaction Reporting Field** setting, move the fields that you want to see in the logs from the **Available** list to the **Selected** list.
13. Click **Finished**.

You have created a rule that sends transactional data about the traffic to external high-speed logging servers. You can now assign the policy to an active subscriber.

Creating a high-speed logging rule for device detection and tethering

You can specify a reporting destination where reports are sent out whenever the subscribers go from a non-tethering state to a tethering state, or vice-versa. Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent. In an enforcement policy, a rule can specify that tethering details are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the high precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.
6. In the **Reporting** setting, specify where to send the tethering detection data:
 - From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs.
 - From the **Format Script** list, select the format script of the report from the **Format Script** list.

Note: The format script is previously configured in **Policy Enforcement > Reporting > Format Script** page.

7. Click **Finished**.

You have created a rule that sends device detection and tethering data about the traffic to external high-speed logging servers.

Session-based reporting format

In an enforcement policy, a rule can send session-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
PEM id	Identifies the reporting module (PEM) and the field value is 23003143.
Version	Indicates the version of the format for backward compatibility.
Timestamp seconds	The time the information was logged (along with the timestamp in milliseconds), specifies seconds using UNIX time format.
Timestamp msec	The time the information was logged (along with the timestamp in seconds), specifies milliseconds using UNIX time format.
Report type	The type of report. Always set to 3 for session-based reporting.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber ID type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
3GPP parameters	The list of 3GPP parameters, which can be imsi, imeisv, tower_id, or username.
Policy ID	The Identification of the policy.
Rule ID	The Identification of the policy rule.
Application ID	A unique number that represents a particular application, and is used for classifying traffic.
Last Sent	The time, in seconds, since the last log entry was sent.
Bytes in	The number of bytes received during this session.
Bytes out	The number of bytes sent during this session.
Concurrent flows	Always 0 (unsupported).
Opened flows	Always 0 (unsupported).
Terminated flows	Always 0 (unsupported).
Total transactions	Always 0 (unsupported).
Successful transactions	Always 0 (unsupported).
Aggregated category duration	Summary of the duration of all flows for the session.
Reason	The reason for sending the record. It can be 0 - reserved, 1 - volume threshold reached, 2- interval time, 3 - subscriber logout, or 4 - inactivity.

Example session-based reporting format

```
Oct 10 17:19:45 172.31.63.64
23003143,1349914925,546879,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914913,5469633,308908379,
0,0,0,0,0,5052,1
Oct 10 17:19:57 172.31.63.64
23003143,1349914937,546661,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914925,5550857,313317479,
0,0,0,0,0,5063,1
Oct 10 17:20:09 172.31.63.64
23003143,1349914949,546676,3,404234567123456,IMSI,linux,f501,
404234567123456,35827001,16394,1349914937,5636605,318053179,
```

0,0,0,0,0,5074,1

Flow-based reporting format

In an enforcement policy, a rule can send flow-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
PEM id	Identifies the reporting module (PEM) and the field value is 2300314.
Version	Indicates the version of the format for backward compatibility.
Timestamp seconds	The time the information was logged in UNIX time format.
Timestamp msec	The msec time value of the timestamp in UNIX time format.
Report type	The type of report; 0 – flow start, 1 – flow interim, 2 – flow end.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber ID type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
Source IP	The IPv4 source address in the IP packet header.
Source port	The source port the subscriber.
Destination IP	The IPv4 destination address in the IP packet header.
Destination IP address	The destination IP of the traffic.
Destination port	The destination port for the traffic.
Protocol	The protocol of the traffic for this flow, TCP or UDP.
Route Domain	The route domain this flow belongs to.
VLAN	The VLAN this flow belongs to.
Application ID	A unique number that represents a particular application in this flow; it is used for classifying traffic.
Urlcat ID	The URL category id that the flow belongs to.
Flow start time seconds	The time, in seconds, the flow started in UNIX time format.
Flow start time msec	The time in milliseconds of the flow start time.
Flow end time seconds	The time the flow ended in UNIX time format.
Flow end time msec	The time in milliseconds of the flow end time.
Transactions count	The count of full transactions seen in the flow.
Bytes in	The number of bytes received during this flow.
Bytes out	The number of bytes sent during this flow.

Example flow-based reporting format

Sep 13 13:48:58 172.31.63.60

```
23003143,1347546777,654398,0,4086007577,E164,2001::10,52784,2001::2,80,6,
67,1347546774,628630,4278124286,4278124286,331,156
Sep 13 13:48:58 172.31.63.60
23003143,1347546777,654398,2,4086007577,E164,2001::10,52784,2001::2,80,6,
67,1347546774,628630,1347546775,382473,547,864
```

Transaction-based reporting format

In an enforcement policy, a rule can send transaction-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
PEM id	Identifies the reporting module (PEM) and the field value is 23003143.
Version	Indicates the version of the format for backward compatibility.
Record type	The type of report; 10 – transactional.
Transaction Number	The sequential number of transaction in this flow (starting from 1).
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber ID type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
Source IP	The IPv4 source address in the IP packet header.
Source port	The source port the subscriber.
Destination IP	The IPv4 destination address in the IP packet header.
Destination port	The destination port for the traffic.
Protocol, TCP/UDP	The protocol of the traffic for this flow, TCP or UDP.
Route Domain ID	The route domain ID of the traffic.
VLAN ID	The VLAN ID of the traffic.
Application/Category ID	A unique number that represents the most relevant application or category that is classified for the transaction.
URL Category ID	A unique number that represents the first (most relevant) URL category that is classified for the transaction.
Transaction Classification result	Reports all classification tokens from the classification engine. <i>Note: The traffic classification result is stored using multiple tokens (8 application/category token identifiers and 4 URL token identifiers) and reported using a CSV format.</i>
Transaction Start, seconds	The transaction timestamp (seconds) in UNIX time format, when an HTTP request is received.
Transaction Start, msec	The transaction timestamp (msec) in UNIX time format when an HTTP request is received.
Transaction Stop, seconds	The transaction timestamp (seconds) in UNIX time format when the corresponding HTTP response is received.
Transaction Stop, msec	The transaction timestamp (msec) in UNIX time format when the corresponding HTTP response is received.

Field	Description
Transaction Upstream Volume, bytes	The number of HTTP request bytes for this transaction.
Transaction Downstream Volume, bytes	The number of HTTP response bytes for this transaction.
Skipped Transactions of this kind	The number of transactional reports skipped within the flow since the last successfully transmission in the flow.
HTTP information:	<p>The HTTP request/response information presented in a CSV format containing the following fields:</p> <ul style="list-style-type: none"> • HTTP Transaction Response Code • HTTP Hostname field truncated (indicates that the Hostname field is truncated due to excessive length) • HTTP Hostname • HTTP User Agent field truncated (indicates that the User Agent field is truncated due to excessive length) • HTTP User Agent • HTTP URI field truncated (indicates that the URI field is truncated due to excessive length) • HTTP URI

Example transaction-based reporting format

```

Jan 15 11:36:27 localhost info tmm[29503]:
23003143,10,1.0.0,1,12341234,IMSI,10.10.10.212,32965,10.10.10.217,80,6,0,311,67,0,
67,16394,0,0,0,0,0,0,0,0,0,1389123382,694,1389123382,697,127,80799103,0,200,
0,10.10.10.217,0,Wget/1.13.4 (linux-gnu),0,/index_long.html
Jan 15 11:36:28 localhost info tmm[29503]:
23003143,10,1.0.0,2,12341234,IMSI,10.10.10.212,32965,10.10.10.217,80,6,0,311,67,0,
67,16394,0,0,0,0,0,0,0,0,0,1389123384,264,1389123384,267,127,80799103,0,200,
0,10.10.10.217,0,Wget/1.13.4 (linux-gnu),0,/index_long.html
Jan 15 11:36:33 localhost info tmm[29503]:
23003143,10,1.0.0,3,12341234,IMSI,10.10.10.212,32965,10.10.10.217,80,6,0,311,67,0,
67,16394,0,0,0,0,0,0,0,0,0,1389123385,572,1389123385,574,127,80799103,0,200,
0,10.10.10.217,0,Wget/1.13.4 (linux-gnu),0,/index_long.html
Jan 15 11:36:33 localhost info tmm[29503]:
23003143,10,1.0.0,4,12341234,IMSI,10.10.10.212,32965,10.10.10.217,80,6,0,311,67,0,
67,16394,0,0,0,0,0,0,0,0,0,1389123387,968,1389123387,970,127,80799103,0,200,
0,10.10.10.217,0,Wget/1.13.4 (linux-gnu),0,/index_long.html
Jan 15 11:36:45 localhost info tmm[29503]:
23003143,10,1.0.0,5,12341234,IMSI,10.10.10.212,32965,10.10.10.217,80,6,0,311,67,0,
67,16394,0,0,0,0,0,0,0,0,0,1389123399,196,1389123399,201,127,80799103,0,200,
0,10.10.10.217,0,Wget/1.13.4 (linux-gnu),0,/index_long.html

```

Device Type and OS-based reporting format

In an enforcement policy, a rule can send Device Type and OS (DTOS)-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
Report id	Identifies the reporting module (PEM) and the field value is 23003143.
Report type	The type of report. Always set to 5 for tethering-based reporting.

Field	Description
Report version	Indicates the version of the format for backward compatibility.
Report timestamp seconds	The time the information was logged (along with the timestamp in milliseconds), specifies seconds using UNIX time format.
Report timestamp millisecs	The time the information was logged (along with the timestamp in seconds), specifies milliseconds using UNIX time format.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
Subscriber IMEISV	The IMEISV value for the subscriber.
Device name	The name of the device obtained from the TacDB.
Device OS	The device OS obtained from the TacDB.
UA-based OS	The OS determined from the user-agent for the sampled flow.
TCP fingerprint-based OS	The OS determined from the TCP fingerprints of the sampled flow.
TCP window size	The Window size from the sampled TCP flow.
Source port	The source port of the sampled flow.
TTL	The time to live (TTL) value of the sampled flow.
TCP header length	The header length of the sampled TCP flow.
TCP window scaling factor	The window scaling factor of the sampled TCP flow.
Record Reason	The reason for sending the record. Set to 14 when tethering is detected, or set to 15 when tethering is not detected.
Operating Systems detected	Not Applicable

Example Device-based reporting format

```

Mar  6 12:07:26 172.31.63.95
23003143,5,1.0.0,1425672465,542,310150123456654,IMSI,012430000098765,Apple-iPhone
 4 model MC603KS,iOS,,,
0,0,0,0,0,14,Linux;Windows;
Mar  6 12:07:56 172.31.63.95
23003143,5,1.0.0,1425672495,542,310150123456654,IMSI,012430000098765,Apple-iPhone
 4 model MC603KS,iOS,,,
0,0,0,0,0,15,
```

Tethering-based reporting format

In an enforcement policy, a rule can send tethering-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
Report id	Identifies the reporting module (PEM) and the field value is 23003143.
Report type	The type of report. Always set to 5 for tethering-based reporting.

Field	Description
Report version	Indicates the version of the format for backward compatibility.
Report timestamp seconds	The time the information was logged (along with the timestamp in milliseconds), specifies seconds using UNIX time format.
Report timestamp millisecs	The time the information was logged (along with the timestamp in seconds), specifies milliseconds using UNIX time format.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Subscriber type	The format of the subscriber ID. It can be E.164, IMSI, NAI, or Private.
Subscriber IMEISV	The IMEISV value for the subscriber.
Device name	The name of the device obtained from the TacDB.
Device OS	The device OS obtained from the TacDB.
UA-based OS	Not Applicable
TCP fingerprint-based OS	Not Applicable
TCP window size	Not Applicable
Source port	Not Applicable
TTL	Not Applicable
TCP header length	Not Applicable
TCP window scaling factor	Not Applicable
Record Reason	The reason for sending the record. Set to 14 when tethering is detected, or set to 15 when tethering is not detected.
Operating Systems detected	Contains the list of operating systems that were detected during the tethering detection sampling interval.

Example: Tethering-based reporting format

```

Mar  6 12:07:26 172.31.63.95
23003143,5,1.0.0,1425672465,542,310150123456654,IMSI,012430000098765,Apple-iPhone
 4 model MC603KS,iOS,,,
0,0,0,0,0,14,Linux;Windows;
Mar  6 12:07:56 172.31.63.95
23003143,5,1.0.0,1425672495,542,310150123456654,IMSI,012430000098765,Apple-iPhone
 4 model MC603KS,iOS,,,
0,0,0,0,0,15,
```


Using Fast L4 for Performance Optimization

Overview: Using Fast L4 for performance enhancement

Policy Enforcement Manager™ (PEM™) provides capability to apply subscriber base policy to multiple subscriber traffic. PEM provides a switch decision between standard or full proxy to BIG-IP® TCP processing per flow, based on preliminary subscriber policy evaluation. You can effectively handle policy actions such as redirect, ICAP adaptation and HTTP header enrichment as the default PEM configuration creates a full proxy listener. Other PEM actions, that do not require full proxy can be performed through Fast L4 basis with BIG-IP TCP support. Some of these PEM actions are quality of service (QoS), bandwidth control, and gate and reporting.

The Performance Layer4 virtual server type uses the Fast L4 profile. The Performance HTTP virtual server type uses the Fast HTTP profile and is designed to speed up certain types of HTTP connections and reduce the number of connections opened to the back-end HTTP servers. This is accomplished by combining features from the TCP, HTTP, and OneConnect™ profiles into a single profile that is optimized for network performance. The Performance HTTP virtual server processes connections on a packet-by-packet basis and buffers only enough data to parse packet headers.

Task summary

Creating a Fast L4 profile in PEM

Disabling a Fast L4 profile in PEM

Creating a Fast L4 profile in PEM

You have to create a listener to handle traffic for policy enforcement. Creating a listener performs preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

The BIG-IP® system specifies the listener that it wants to use for a flow in situation where it decides the optimal listener. You can add an option to specify the listener name in the configuration of the PEM profile.

1. On the Main tab, click **Policy Enforcement > Profiles > Policy Enforcement**.
The Policy Enforcement screen opens.
2. Click **Create**.
The New PEM Profile screen opens.
3. In the **Name** field, type a unique name of the policy enforcement profile.
4. In the Optimization area, for the **Connection Optimization** setting, **Enabled** is the default selection. This indicates that the fast PEM optimization is enabled. To change the setting, select **Custom** check box and select **Disabled**.
5. From the **Fast L4 Virtual Server** list, select the Fast L4 virtual server previously configured. The Fast L4 virtual server is a server with an HTTP profile, protocol TCP, and protocol profile Fast L4.
6. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

7. Click **Finished**.

The Policy Enforcement Manager creates a policy enforcement profile with Fast L4.

You have created a Fast L4 profile.

Disabling a Fast L4 profile in PEM

You have to create a listener to handle traffic for policy enforcement. Creating a listener performs preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

The BIG-IP® system specifies the listener that it wants to use for a flow in situation where it decides the optimal listener. You can add an option to specify the listener name in the configuration of the PEM™ profile.

1. On the Main tab, click **Policy Enforcement > Profiles > Policy Enforcement**.
The Policy Enforcement screen opens.
2. From the PEM profile list, select a Fast L4 profile that you have created, or click **Create** to create a new Fast L4 profile.
3. In the **Name** field, type a unique name of the policy enforcement profile.
4. In the Optimization area, for the **Connection Optimization** setting, **Enabled** is the default selection. This indicates that the fast PEM optimization is enabled. To change the setting, select **Custom** check box and select **Disabled**.
5. From the **Fast L4 Virtual Server** list, select the Fast L4 virtual server previously configured. The Fast L4 virtual server is a server with an HTTP profile, protocol TCP, and protocol profile Fast L4.
6. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

7. Click **Finished**.

The Policy Enforcement Manager creates a policy enforcement profile with Fast L4.

You have disabled Fast L4 on a PEM profile.

Reporting Quality of Experience and Video Usage

Overview: Reporting quality of experience and video usage

In Policy Enforcement Manager™ (PEM), you can gather report on subscriber video traffic consumption across different devices (phone, tablet, PC or TV). This information helps to analyze user quality of experience (QoE).

After a PEM and QoE profile is enabled on a virtual, the QoE module detects video flows and gathers data. When the QoE data is passed to PEM, a subscriber QoE report is generated based on the corresponding policy settings. The reports can be sent over syslog (HSL) or IPFIX.

PEM has visibility in to the following components:

- Subscriber
- The media traffic
- Control for video optimization
- The point of control for TCP optimization.
- Custom policy provisioning using iRules®

Task summary

Creating a publisher

Adding Quality of Experience profile to the virtual server

Configuring QoE Reporting

Creating a publisher

Before you create a publisher, you have to create a HSL pool that needs to be associated to a destination. Ensure that at least one destination associated with a pool of remote log servers exists on the BIG-IP® system.

Create a publisher to specify where the BIG-IP system sends log messages for specific resources.

1. On the Main tab, click **System > Logs > Configuration > Log Publishers**.
The Log Publishers screen opens.
2. Click **Create**.
3. In the **Name** field, type a unique, identifiable name for this publisher.
4. For the **Destinations** setting, select a destination from the **Available** list, and click << to move the destination to the **Selected** list.

***Note:** If you are using a formatted destination, select the destination that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

5. Click **Finished**.

Adding Quality of Experience profile to the virtual server

Before creating a video Quality of Experience (QoE) virtual server, you need to have created and configured a video QoE profile.

You can assign video QoE profile to a virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. From the **HTTP Profile** list, select **http**.
4. In the Resources area, for the **iRules** setting, from the **Available** list, select the name of the QoE iRule that you want to assign, and move the name into the **Enabled** list.
5. In the Resources area of the screen, from the **Default Pool** list, select the relevant pool name.
6. Click **Finished**.
7. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
8. Click the name of the virtual server you want to modify.
9. From the **Configuration** list, select **Advanced**.
10. From the **QoE** list, select a QoE profile to attach to the virtual server.
11. Click **Update** to save the changes.

This assigns the video QoE profile and iRules to the virtual server.

Configuring QoE Reporting

Before you can enable the Quality of Experience (QoE) attribute at a policy level, you have to enable QoE in the profile section.

In an enforcement policy, a media-quality QoE report action can be added, that can be added with other reporting actions in the same rule.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for

the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **QoE Reporting** list, select **Enabled**.
8. In the **QoE Destination** setting, from the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs and select the format script of the report from the **Format Script** list.

***Note:** If you are using a formatted destination, select the publisher that matches your log servers, such as Remote Syslog, Splunk, or ArcSight.*

9. Click **Finished**.

You have created an enforcement policy with QoE report action.

QoE-based reporting format

In an enforcement policy, a rule can send QoE-based information about traffic that matches certain criteria to an external high-speed logging (HSL) server. The logs include the following comma-separated values in the order listed.

Field	Description
Report id	Identifies the reporting module (PEM) and the field value is 23003143.
Subscriber ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
Source IP	The IPv4 source address in the IP packet header.
Source Transport Port	The source (L4) port.
Destination IP	The IPv4 destination address in the IP packet header.
Destination Transport Port	The IPv4 destination address in the IP packet header.
Protocol Identifier	The IP Protocol field.
Media Type	Different types of media, for example, MP4.
URL X SessionId	The ID used to associate different segments of a whole video or audio.
Width Height	The resolution of the video.
Bit Rate	The number of bits that are conveyed or processed per unit of time.
Frame Rate	The frequency (rate) at which an imaging device produces unique consecutive images called frames.
Duration	The length of time of the media.
Watched	It is the length of time that the video has been watched.
Mos	It is the value ranging 1 to 5, that evaluates the user-experience.

Example QoE-based reporting format

```
Apr 30 14:30:14 slot2/sush vic 172 info tmm[4243]:  
23003143,6,1.0.0,1430429414,610,404234567123456,IMSI,10.1.1.11,37112,11.1.1.100,80,6,2426616,0,320x240,FLV,,1,0,5  
Apr 30 14:30:46 slot2/sush vic 172 info tmm[4243]:  
23003143,6,1.0.0,1430429446,88,404234567123456,IMSI,10.1.1.11,37113,11.1.1.100,80,6,164771,0,480x320,MP4,,70,0,4
```

Performing Radius Authentication and Accounting

Overview: Performing RADIUS authentication and accounting

In Policy Enforcement Manager™, the RADIUS client has the ability to initiate RADIUS authentication for a subscriber. You can configure the virtual servers that are used to request for authentication of DHCPv4 and DHCPv6 discovered subscribers. The subscriber authentication may be triggered by subscriber discovery based on other means, such as obtaining RADIUS accounting messages. The ability to generate accounting messages helps to track subscriber usage as a RADIUS client.

RADIUS authentication is initiated when PEM receives messages, showing that the subscribers are attempting to connect to the network. The two factors of initiation are:

- The start of DHCP exchange showing that the subscriber attempts to obtain an IP address (fixed line deployments).
- When the RADIUS accounting start message indicates that the subscriber has passed through the initial phase of access but still needs authentication.

Task summary

Creating a RADIUS AAA profile for policy enforcement

Creating a listener for RADIUS AAA Virtual

Creating policy rule for RADIUS accounting reports

Creating a RADIUS AAA profile for policy enforcement

Create a RADIUS profile, which contains the shared secret of the RADIUS server, the transaction timeout, password, and retransmission timeout details, for configuring the RADIUS authentication profile settings.

1. On the Main tab, click **Local Traffic > Profiles > Policy Enforcement > RADIUS AAA**.
2. Click **Create**.
The New Radius Profile screen opens.
3. In the **Name** field, type a unique name for the profile.
4. In the **Description** field, type a descriptive text that identifies the profile.
5. From the **Parent Profile** list, select the default **radiusaaa** profile.
6. Select the **Custom** check box.
7. For the **Secret** setting, select the **Custom** check box to enable this option. Type the shared secret of the RADIUS server used for authentication.
8. For the **Password** setting, select the **Custom** check box to enable this option. Type the password of the RADIUS AAA profile for RADIUS server authentication.
9. For the **Transaction Timeout** setting, select the **Custom** check box to enable this option. Type the number, in seconds, of the time taken for server to respond.
10. For the **Retransmission Timeout** setting, select the **Custom** check box to enable this option. Type the number of seconds to wait before resending authentication or accounting messages to the RADIUS server.

The RADIUS profile that you created can be chosen from the RADIUS profile in **Local Traffic > Virtual Servers > Virtual Server List > New Virtual Server >** , depending on the virtual server IP address type.

Creating a listener for RADIUS AAA Virtual

You can create new RADIUS AAA virtuals to authenticate or send accounting information about the subscriber to the RADIUS server.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. From the Authentication Virtuals area, click **Add**.
The New RADIUS AAA Virtual screen opens.
3. In the **Name** field, type a unique name for the RADIUS AAA virtual.
4. In the **Description** field, type a description of the listener.
5. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
6. From the **Mode** list, select the **Authentication** or **Accounting** to specify the type of RADIUS virtual you are creating.
7. For the **Secret** setting, select the **Custom** check box to enable this option. Type the shared secret of the RADIUS server used for authentication or accounting.
8. For the **Password** setting, select the **Custom** check box to enable this option. Type the password of the RADIUS AAA profile for RADIUS server authentication.
9. For the **Pool Member Configuration** setting, add the RADIUS AAA virtual servers that are to be members of the pool. Type the **Member IP Address** and **Port** number, then click **Add**.
You can use port 1812 for RADIUS authentication and port 1813 for RADIUS accounting.
10. Click **Finished**.
The Policy Enforcement Manager creates a RADIUS AAA virtual server, and displays in the authentication virtuals list.

When you create a RADIUS AAA virtual for a subscriber, the Policy Enforcement Manager™ initiates RADIUS authentication or sends accounting information, for that subscriber. A RADIUS AAA profile is also created and is assigned to the virtual server automatically.

Creating policy rule for RADIUS accounting reports

Policy Enforcement Manager™ (PEM™) allows you to specify a RADIUS internal virtual server as a reporting destination. The reporting thresholds are optional if RADIUS destination is selected.

Note: Only one reporting destination can be specified in a given rule.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.

5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Usage Reporting** list, select **Enabled**.
8. From the **Report Granularity** list, select from one the the granular reporting options:

Option	Description
Session	Select Session to log details about subscribers and application sessions.
Flow	Select Flow , for more granular reporting of every TCP connection.
Transaction	select Transaction , for more granular reporting of every HTTP transaction.

9. If you select **Session** or **Flow**, in the **Volume Threshold** setting, specify in octets, the threshold to send RADIUS reporting records. You can send reporting data from uplink traffic, to downlink traffic and the total traffic volume before logging the information.
10. If you select **Transaction**, in the **Additional HTTP Information** setting, specify in bytes, the HTTP **Hostname**, the HTTP **User Agent** and the HTTP **URI**.
11. In the **Destination** setting, Select the **RADIUS Accounting** option from the destination.
12. From the **RADIUS AAA Virtual** list, select the RADIUS AAA virtual that you created earlier.
13. Click **Finished**.

You have created a RADIUS internal virtual server as a reporting destination.

Configuring Subscriber Discovery based on DHCP

Overview: Configuring subscriber discovery based on DHCP

The Policy Enforcement Manager™ uses DHCP to discover subscribers. The DHCP consists of two components, which includes a protocol for delivering host-specific parameters from a DHCP server to a host, and the ability to allocate network addresses to hosts. The BIG-IP® system processes the DHCP traffic between subscribers and DHCP server and extracts the subscriber's identity and other information that is important for subscriber handling.

The BIG-IP DHCP module has two functional modes:

- Relay mode: The DHCP-Relay agent handles the DHCP traffic from the subscriber, modifies it as required, and relays it to the DHCP server according to the configuration.
- Forward or pass-through mode: The DHCP module does not relay the messages or modify the message in this mode.

In both modes, the DHCP module snoops the DHCP packets, parses relay-agent options and the allocated IP address, and then extracts session information. The relay-agent options are option 82 for DHCPv4 and options 37 and 38 for DHCPv6.

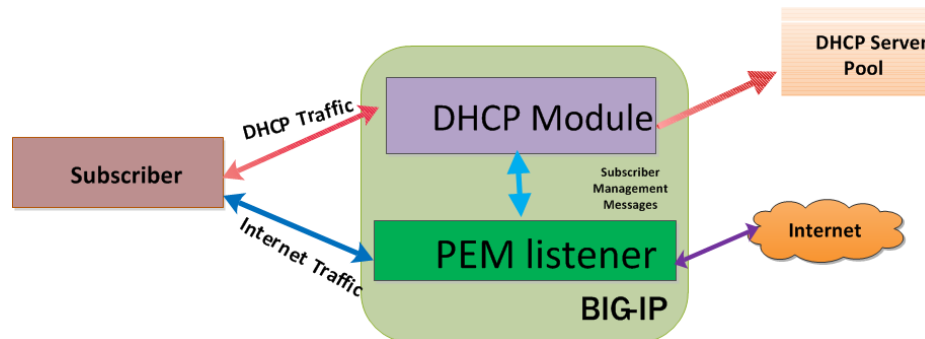


Figure 2: Subscriber Discovery through DHCP

The DHCP module monitors the clients DHCP traffic after the initial IP allocation and snoops for DHCP lease renewal packets, releasing of the IP address, and reconfiguring requests. This determines when the BIG-IP system can safely delete the session.

Task summary

- Creating a listener for DHCPv4 discovery virtual*
- Creating a DHCPv4 profile for policy enforcement*
- Creating a listener for DHCPv6 discovery virtual*
- Creating a DHCPv6 profile for policy enforcement*
- Creating a listener for RADIUS subscriber discovery*

Creating a listener for DHCPv4 discovery virtual

You can use DHCP to discover subscribers in order to handle traffic for policy enforcement. For subscribers discovered through DHCP, an identifier comprises of relay agent information option (option 82) and MAC address, as configured in the corresponding DHCP profile.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Select **DHCPv4** from the profiles list, and click **Add**.
The New DHCPv4 Discovery Virtual screen opens.
3. In the **Name** field, type a unique name for the listener.
4. In the **Description** field, type a description of the listener.
5. For the **Source** setting, type the IP address or network from which the virtual server will accept traffic.
6. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For the **DHCP Mode** setting, select **Relay** or **Forward** to specify the mode in which the DHCP client requests are sent.
10. For the **Pool Member Configuration** setting, add the DHCP virtual servers that are to be members of the pool. Type the **Member IP Address** and **Port** number, then click **Add**.
11. From the **Subscriber Discovery** list, select **Enabled**. Then, for the **Subscriber ID Format** setting, select the format you want to implement.

Format	Description
MAC Address	Uses the subscriber ID as the MAC address through which the subscriber ID goes through.
Relay Agent Option: Suboption ID 1	Uses the relay agent first option suboption ID.
Relay Agent Option: Suboption ID 1 + <Separator> + Suboption ID 2	Uses the relay agent first and second suboption IDs.
MAC Address + <Separator> + Relay Agent Option: Suboption ID 1	Uses the MAC Address and the relay agent first suboption ID.
MAC Address + <Separator> + Relay Agent Option: Suboption ID 1 <Separator> + Suboption ID 2	Uses the relay agent first option suboption ID.
TCL Expression	Uses the TCL expression to format the subscriber ID.

12. From the **Authentication Settings** list, select **Enabled**. Then, select the virtual server name from the **Authentication Virtual** list. Select the **User Name Format** you want to implement.
The **User Name Format** has the same options as the **Subscriber ID Format**, in the Subscriber Discovery setting.
13. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a new DHCPv4 discover virtual, the Policy Enforcement Manager™ also creates a corresponding DHCPv4 profile.

Creating a DHCPv4 profile for policy enforcement

You can create a DHCP profile when you want to configure the DHCP virtual to use Relay mode or Pass-through mode.

1. On the Main tab, click **Local Traffic > Profiles > Services > DHCPv4**.
2. Click **Create**.
The New DHCPv4 Profile screen opens.
3. In the **Description** field, type a descriptive text that identifies the profile.
4. From the **Parent Profile** list, select the default **dhcpv4** profile.
5. Select the **Custom** check box.
6. In the Protocol and Proxy Settings Features area, make a selection from the **DHCP Mode** list.

Option	Description
Relay	When in relay mode, a virtual server relays Dynamic Host Control Protocol (DHCP) client requests and applies unicast IP addresses as the relayed message destination.
Forward	When in forward mode, a virtual server forwards Dynamic Host Control Protocol (DHCP), and does not modify, client requests for an IP address to one or more DHCP servers.
7. For the **Idle Timeout** setting, type the number of seconds that a BIG-IP DHCP connection is idle before the connection is eligible for deletion.
8. For the **Max Hops** setting, select the **Custom** check box to enable this option. Type the maximum expected number of relay agents that the messages should pass through, before reaching the DHCPv4 server.
9. For the **Default TTL** setting, select the **Custom** check box to enable this option. Type the time to live (TTL) value that you want to set for each outgoing DHCP packet.
10. For the **Default Lease Time** setting, select the **Custom** check box to enable this option. Type the time, in seconds, of the default value of the DHCPv4 lease time.
11. For the **TTL Decrement Amount** setting, select the **Custom** check box to enable this option. Type the amount that the DHCP virtual will use to decrement the TTL for each outgoing DHCP packet.
12. For the **Transaction Timeout** setting, select the **Custom** check box to enable this option. Type the number of seconds, taken to internally process the messages.
13. For the **Insert Relay Agent ID (Option 82)** setting, select the **Custom** check box to enable this option if you want the DHCP module to insert option 82.
14. For the **Remove Relay Agent ID From Client Messages** setting, select the **Custom** check box to enable this option and if you want the DHCP relay agent to remove option 82 from the server to client traffic.

15. From the **Subscriber Discovery** list, select **Enabled**. Then, for the **Subscriber ID Format** setting, select the format you want to implement.

Format	Description
MAC Address	Uses the subscriber ID as the MAC address through which the subscriber ID goes through.
Relay Agent Option: Suboption ID 1	Uses the relay agent first option suboption ID.
Relay Agent Option: Suboption ID 1 + <Separator> + Suboption ID 2	Uses the relay agent first and second suboption IDs.
MAC Address + <Separator> + Relay Agent Option: Suboption ID 1	Uses the MAC Address and the relay agent first suboption ID.
MAC Address + <Separator> + Relay Agent Option: Suboption ID 1 <Separator> + Suboption ID 2	Uses the relay agent first option suboption ID.
TCL Expression	Uses the TCL expression to format the subscriber ID.

16. From the **Authentication Settings** list, select **Enabled**. Then, select the virtual server name from the **Authentication Virtual** list. Select the **User Name Format** you want to implement.
The **User Name Format** has the same options as the **Subscriber ID Format**, in the Subscriber Discovery setting.
17. Click **Finished**.

The DHCPv4 profile that you created can be chosen from the DHCPv4 profiles in **Local Traffic > Virtual Servers > Virtual Server List > New Virtual Server >** , only if you choose DHCP as a virtual type.

Creating a listener for DHCPv6 discovery virtual

You can use DHCPv6 to discover subscribers in order to handle traffic for policy enforcement. For each subscriber discovered through DHCPv6, an identifier comprises of remote-id, subscriber-id options (options 37 and 38) and MAC address, as configured in the corresponding DHCPv6 profile.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Select **DHCPv6** from the profiles list, and click **Add**.
The New DHCPv6 Discovery Virtual screen opens.
3. In the **Name** field, type a unique name for the listener.
4. In the **Description** field, type a description of the listener.
5. For the **Source** setting, type the IP address or network from which the virtual server will accept traffic.
6. In the **Destination Address** field, type the IP address of the virtual server. For example, `ff01::0020/64` or `2001:ed8:77b5:2:10:10:100:42/64`.

Tip: For DHCPv6 discovery virtual, the source and destination should be any (::/0).

The system will create a virtual server using the address or network you specify.

7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.

8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For the **DHCP Mode** setting, select **Relay** or **Forward** to specify the mode in which the DHCP client requests are sent.
10. For the **Pool Member Configuration** setting, add the DHCP virtual servers that are to be members of the pool. Type the **Member IP Address** and **Port** number, then click **Add**.
11. From the **Subscriber Discovery** list, select **Enabled**. Then, for the **Subscriber ID Format** setting, select the format you want to implement.

Format	Description
MAC Address	Uses the subscriber ID as the MAC address through which the subscriber ID goes through.
MAC Address + <Separator> + Option 37	Uses the MAC address and the remote ID relay agent option.
MAC Address + <Separator>+ Option 37 <Separator> + Option 38	Uses the MAC address, the remote ID relay agent option and the subscriber ID option.
MAC Address + <Separator> + Option 38	Uses the MAC address and the subscriber ID option.
Option 37	Uses the remote ID relay agent option.
Option 37 <Separator> + Option 38:	Uses the remote ID relay agent option and the subscriber ID option.
Option 38	Uses the subscriber ID option.
TCL Expression	Uses the TCL expression to format the subscriber ID.

12. From the **Authentication Settings** list, select **Enabled**. Then, select the virtual server name from the **Authentication Virtual** list. Select the **User Name Format** you want to implement.
The **User Name Format** has the same options as the **Subscriber ID Format**, in the Subscriber Discovery setting.

When you create a new DHCPv6 discover virtual, the Policy Enforcement Manager™ also creates a corresponding DHCP profile.

Creating a DHCPv6 profile for policy enforcement

You can create a DHCP profile when you want to configure the DHCP virtual to use Relay mode or Pass-through mode.

1. On the Main tab, click **Local Traffic > Profiles > Services > DHCPv6**.
2. In the **Description** field, type a descriptive text that identifies the profile.
3. From the **Parent Profile** list, select the default **dhcpv6** profile.
4. Select the **Custom** check box.
5. In the Protocol and Proxy Settings Features area, make a selection from the **DHCP Mode** list.

Option	Description
Relay	When in relay mode, a virtual server relays Dynamic Host Control Protocol (DHCP) client requests and applies unicast IP addresses as the relayed message destination.

Option	Description
Forward	When in forward mode, a virtual server forwards Dynamic Host Control Protocol (DHCP), and does not modify, client requests for an IP address to one or more DHCP servers.

6. For the **Idle Timeout** setting, type the number of seconds that a BIG-IP DHCP connection is idle before the connection is eligible for deletion.
7. For the **Max Hops** setting, select the **Custom** check box to enable this option. Type the maximum expected number of relay agents that the messages should pass through, before reaching the DHCPv4 server.
8. For the **Default Lease Time** setting, select the **Custom** check box to enable this option. Type the time, in seconds, of the default value of the DHCPv4 lease time.
9. For the **Transaction Timeout** setting, select the **Custom** check box to enable this option. Type the number of seconds, taken to internally process the messages.
10. For the **Insert Remote ID (Option 37)** setting, select the **Custom** check box to enable this option if you want the DHCP module to insert option 37.
11. For the **Insert Remote ID (Option 37)** setting, select the **Custom** check box to enable this option if you want the DHCP module to insert option 38.
12. For the **Remove Subscriber Agent ID From Client Messages** setting, select the **Custom** check box to enable this option and if you want the DHCP relay agent to remove option 37 from the server to client traffic.
13. For the **Remove Relay Agent ID From Client Messages** setting, select the **Custom** check box to enable this option and if you want the DHCP module to remove option 38 from the server to client traffic.
14. From the **Subscriber Discovery** list, select **Enabled**. Then, for the **Subscriber ID Format** setting, select the format you want to implement.

Format	Description
MAC Address	Uses the subscriber ID as the MAC address through which the subscriber ID goes through.
MAC Address + <Separator> + Option 37	Uses the MAC address and the remote ID relay agent option.
MAC Address + <Separator>+ Option 37 <Separator> + Option 38	Uses the MAC address, the remote ID relay agent option and the subscriber ID option.
MAC Address + <Separator> + Option 38	Uses the MAC address and the subscriber ID option.
Option 37	Uses the remote ID relay agent option.
Option 37 <Separator> + Option 38:	Uses the remote ID relay agent option and the subscriber ID option.
Option 38	Uses the subscriber ID option.
TCL Expression	Uses the TCL expression to format the subscriber ID.

15. From the **Authentication Settings** list, select **Enabled**. Then, select the virtual server name from the **Authentication Virtual** list. Select the **User Name Format** you want to implement.
The **User Name Format** has the same options as the **Subscriber ID Format**, in the Subscriber Discovery setting.
16. Click **Finished**.

The DHCPv6 profile that you created can be chosen from the DHCPv6 profiles in **Local Traffic > Virtual Servers > Virtual Server List > New Virtual Server >** , only if you choose DHCP as a virtual type.

Creating a listener for RADIUS subscriber discovery

You can create listeners that specify the RADIUS discovery virtual for extracting subscriber information from the RADIUS packets. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. From the Subscriber Discovery Virtuals area, select **RADIUS**, and click **Add**.
The New RADIUS Discovery Virtual screen opens.
3. In the **Name** field, type a unique name for the RADIUS discovery virtual.
4. In the **Description** field, type a description of the listener.
5. For the **Source** setting, type the IP address or network from which the virtual server will accept traffic.
6. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
9. For the **Pool Member Configuration** setting, add the RADIUS discovery virtual servers that are to be members of the pool. Type the **Member IP Address** and **Port** number, then click **Add**.
10. Click **Finished**.
The Policy Enforcement Manager creates a RADIUS virtual server, and displays in the subscriber discovery list.

When you create a RADIUS discovery virtual for a subscriber, the Policy Enforcement Manager™ creates a corresponding profile (**Policy Enforcement > Listeners > Control Virtual Servers**).

Usage Monitoring Over a Gx Interface

Overview: Usage monitoring over a Gx and Sd interface

In Policy Enforcement Manager™, you can create a rule within an enforcement policy that tells the system to send aggregated usage data concerning individual subscribers to a Policy and Charging Rules Function (PCRF). The rule specifies what type of traffic you are interested in, and one of the actions the system can take with the data collected is to send it for processing over a Gx interface to a PCRF.

The system sends the data in the standard Gx format. The report granularity must be set to session for Gx reporting to be available. The PCRF determines the policies for each subscriber, whether or not reporting is enabled, and how often to send the data and monitoring key that identifies the type of traffic PCRF wants to get usage for.

Similarly, the Traffic Detection function (TDF) functionality performs applications traffic detection and reporting of detected application by using TDF application identifier. The Policy Enforcement Manager™ supports the Sd interface which is used by PCRF to talk to TDF. This provides the ability to apply the detection, enforce actions and apply charging parameters for the specified application traffic.

Note: *Gx and Sd are mutually exclusive.*

For example, a rule might collect session-based information about all traffic destined to a particular IP address. The BIG-IP® system communicates with the PCRF and sends information about the subscribers for whom reporting is enabled. You establish the connection to the PCRF by creating a listener with Gx interface enabled.

Task summary

Creating a listener for subscriber discovery and policy provisioning

Creating a rule for usage monitoring

Creating a listener for subscriber discovery and policy provisioning

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. In the Policy Provisioning and Online Charging Virtuals area, click **Add**.
The New Configure Diameter Endpoint Provisioning and Online Charging screen opens.
3. In the **Name** field, type a unique name for the listener.
4. In the **Description** field, type a description of the listener.
5. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
6. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** or **Sd** from the **Supported Apps** options.

7. In the **Product Name** field, type the product name which is used to communicate with the PCRF.
8. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
9. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
10. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
11. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
12. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
13. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
14. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
15. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
16. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Creating a rule for usage monitoring

In an enforcement policy, a rule can specify that usage monitoring statistics concerning traffic affected by the rule are sent to a Gx interface.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Usage Reporting** list, select **Enabled**.
8. From the **Report Granularity** list, select **Session**.
You can send only session-based reporting data over the Gx interface.
9. For the **Volume Threshold** setting, specify, in octets, the amount of data to receive from the client, send to the client, and the total traffic volume before logging the information.
10. For the **Destination** setting, specify these values:
 - a) For **Gx** or **Sd**, select **Enabled**.
 - b) In the **Gx Monitoring Key** or **Sd Monitoring Key** field (as selected), type a string to use for usage monitoring of the service data that the enforcement policy rule or dynamic policy and charging control (PCC) rule controls.
11. Click **Finished**.

You have created a rule that sends data about the traffic to the Gx interface in the standard Gx format.

Configuring Global Application Policies with Bandwidth Control

Overview: Global Application Policies with Bandwidth Control

You can use bandwidth controllers with Policy Enforcement Manager™ to restrict bandwidth usage per subscriber, group of subscribers, per application, and so on. This implementation uses PEM for global application control to limit overall bandwidth for all P2P (or other application) traffic. For example:




Figure 3: Diagram of bandwidth usage per subscriber

In the figure, three subscribers have individual policies that allow P2P bandwidths of up to 4 Mbps each. The maximum bandwidth for all P2P traffic is limited to 10 Mbps (specified as the maximum rate in a static bandwidth controller). If all were sending P2P traffic, they would all get less bandwidth if you apply a global enforcement policy that enforces bandwidth control.

For this implementation, you create the bandwidth controller and the enforcement policy on the BIG-IP® system. In the enforcement policy, a rule applies bandwidth control to P2P traffic. From the listener, you apply the policy globally to all traffic.

Task Summary

Creating VLANs

Creating a static bandwidth control policy

Creating an enforcement policy

Creating a rule for bandwidth control

Creating a listener: example

Creating VLANs

VLANs represent a collection of hosts that can share network resources, regardless of their physical location on the network. For Policy Enforcement Manager™, you typically create VLANs for the subscriber traffic coming in to the BIG-IP® system, for traffic going out to the network, and if using w-steering with service chains, you need two VLANs for each value added service to be fully transparent.

1. On the Main tab, click **Network > VLANs**.
The VLAN List screen opens.
2. Click **Create**.
The New VLAN screen opens.
3. In the **Name** field, type a unique name for the VLAN.
4. In the **Tag** field, type a numeric tag, between 1-4094, for the VLAN, or leave the field blank if you want the BIG-IP system to automatically assign a VLAN tag.
The VLAN tag identifies the traffic from hosts in the associated VLAN.
5. For the **Interfaces** setting:
 - a) From the **Interface** list, select an interface number or trunk name.
 - b) From the **Tagging** list, select **Tagged** or **Untagged**.
Select **Tagged** when you want traffic for that interface to be tagged with a VLAN ID.
 - c) If you specified a numeric value for the **Customer Tag** setting and from the **Tagging** list you selected **Tagged**, then from the **Tag Mode** list, select a value.
 - d) Click **Add**.
 - e) Repeat these steps for each interface or trunk that you want to assign to the VLAN.
6. From the **Configuration** list, select **Advanced**.
7. If you want the system to verify that the return route to an initial packet is the same VLAN from which the packet originated, select the **Source Check** check box.
8. In the **MTU** field, retain the default number of bytes (**1500**).
9. If you want to base redundant-system failover on VLAN-related events, select the **Fail-safe** box.
10. From the **CMP Hash** list, select the appropriate value depending on the location of the VLAN in the system:
 - On the VLAN coming in to the BIG-IP system (often called *internal*), select **Source Address**.
 - On VLANs going out (often called *external*), leave the value set to **Default**.
 - For traffic returning to the BIG-IP from the Internet, select **Destination Address**.
 - If using w-steering for value-added services, on the VLAN coming back to the BIG-IP system, select **Source Address**.
11. Click **Finished**.
The screen refreshes, and displays the new VLAN in the list.

Create as many VLANs as needed for your configuration.

Creating a static bandwidth control policy

You can create a static bandwidth control policy to limit the bandwidth that traffic uses on the BIG-IP® system.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for the resource you are managing.
The number must be in the range from 1 Mbps to 320 Gbps. This value is the amount of bandwidth available to all the connections going through this static policy.
5. Click **Finished**.

For the bandwidth control policy to take effect, you must apply the policy to traffic, using a virtual server, packet filter, or route domain.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.

***Tip:** When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.*

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

***Important:** The system performance is significantly affected, depending on complexity of the classification and the type of policy action.*

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Creating a rule for bandwidth control

If you want to use rate control, you need to have already created a bandwidth controller.

You can create a rule that provides bandwidth control. For example, the bandwidth controller might limit the total amount of bandwidth that can be used by application traffic, such as P2P.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it*

has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. In the **Classification** setting, filter the application traffic to which you want to apply bandwidth control.
 - a) For **Match Criteria**, select **Match**.
 - b) For **Category**, select **P2P** (or other application traffic you want to limit on the network).
 - c) Click **Add**.
7. In the Forwarding area, ensure that **Gate Status** is set to **Enabled**.
8. In the **Rate Control** setting, for **Bandwidth Controller**, select the name of the bandwidth controller that you created to limit P2P (or other application) traffic.
9. Click **Finished**.

You have created a rule to restrict the total bandwidth usage for all P2P traffic to the **Maximum Rate** specified in the static bandwidth control policy.

The enforcement policy needs to be associated with the virtual servers required for PEM. You can do this by creating a listener (recommended), or you can edit the virtual servers to specify the enforcement policy as a global policy, and enable classification.

Creating a listener: example

You create a listener to complete the preliminary setup on the BIG-IP® system; in this case, to apply bandwidth management as a global enforcement policy.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Add**.
The New Virtual Group screen opens.
3. In the **Name** field, type a unique name for the listener.
4. For the **Source** setting, type the IP address or network from which the virtual server will accept traffic.
5. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

6. For the **Service Port** setting, type or select the service port for the virtual server.
7. Subscriber provisioning using RADIUS is enabled by default. If your system is using RADIUS for snooping subscriber identity, you need to specify VLANs and tunnels. If you are not using RADIUS, you need to disable it.
 - For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
 - If you do not want to use RADIUS, from the **Subscriber Identity Collection** list, select **Disabled**.

8. In the Policy Provisioning area, for **Global Policy**, move the enforcement policy you created for bandwidth control to **High Precedence**.
The system applies the policy with bandwidth control to all traffic.
9. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. All traffic classified as P2P traffic is limited to the **Maximum Rate** specified in the static bandwidth control policy. Once the maximum rate is reached, no additional P2P traffic is allowed on the network.

Managing Traffic with Bandwidth Controllers

Overview: Bandwidth control management

Fine-grained bandwidth control is essential to service providers, large enterprises, and remote access services (RAS) solutions. Bandwidth controllers on the BIG-IP® system can scale easily, work well in a distributed environment, and are easy to configure for various networks. Depending on the type of policy you configure, you can use bandwidth controllers to apply specified rate enforcement to traffic flows or mark traffic that exceeds limits.

Bandwidth control policies can be static or dynamic. Through the user interface (browser or `tmsm` command-line utility), when you apply a bandwidth control policy to a virtual server, packet filter, or route domain, you can apply only one policy at a time, and that is a static policy. Using iRules®, you can combine static and dynamic bandwidth control policies up to eight policies on a connection, but only one of the eight policies can be a dynamic policy. A packet is transmitted only when all the attached policies allow it. The system as a whole supports a maximum of 1024 policies.

Important: *Applying a bandwidth controller policy to a route domain affects all traffic transmitted by the BIG-IP system to VLANs in the route domain, including health monitors and DNS queries.*

Note: *Only static bandwidth control policies support SNMP queries.*

Bandwidth controllers vs. rate shaping

Bandwidth controller is the updated version of rate shaping on the BIG-IP® system. These features are mutually exclusive. You can configure and use either rate shaping or bandwidth controllers, but not both. Bandwidth controllers include distributed control, subscriber fairness, and support for a maximum rate of 320 Gbps. Rate shaping is hierarchical and supports minimum bandwidth (committed information rate), priority, and flow fairness.

About static bandwidth control policies

A *static* bandwidth control policy controls the aggregate rate for a group of applications or a network path. It enforces the total amount of bandwidth that can be used, specified as the maximum rate of the resource you are managing. The rate can be the total bandwidth of the BIG-IP® device, or it might be a group of traffic flows.

Task summary for creating a static bandwidth control policy

This procedure includes the steps for assigning a static bandwidth control policy to traffic, using a virtual server. Alternatively, you can assign a static bandwidth control policy to a packet filter or a route domain.

Task list

Creating a static bandwidth control policy

Adding a static bandwidth control policy to a virtual server

Creating a static bandwidth control policy

You can create a static bandwidth control policy to limit the bandwidth that traffic uses on the BIG-IP® system.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for the resource you are managing.
The number must be in the range from 1 Mbps to 320 Gbps. This value is the amount of bandwidth available to all the connections going through this static policy.
5. Click **Finished**.

For the bandwidth control policy to take effect, you must apply the policy to traffic, using a virtual server, packet filter, or route domain.

Adding a static bandwidth control policy to a virtual server

Adding a static bandwidth control policy to a virtual server is one way to apply the policy to traffic. Alternatively, you can add the bandwidth control policy to a packet filter or a route domain.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. From the **Configuration** list, select **Advanced**.
4. From the **Bandwidth Controller** list, select a bandwidth control policy.
Only static bandwidth control policies are available in this list.
5. Click **Update** to save the changes.

The BIG-IP® system now applies rate enforcement to the traffic intercepted by this virtual server, according to the static bandwidth policy you selected. A static bandwidth policy associated with a virtual server applies only to client-initiated flows, and not to bandwidth for traffic flowing toward the client.

About dynamic bandwidth control policies

You can create dynamic bandwidth control policies to restrict bandwidth usage per subscriber or group of subscribers, per application, per network egress link, or any combination of these. A *dynamic* bandwidth control policy provides fairness on traffic flows, according to configurable parameters, within an upper bandwidth limit. The BIG-IP® system activates the dynamic bandwidth control policy for each user only when the user participates. When you create a dynamic bandwidth control policy, it acts as a policy in waiting, until the system detects egress traffic that matches the traffic you want to control and creates an instance of the policy. At that moment, the system applies the bandwidth control policy limits, as specified. No bandwidth control occurs until the system detects traffic and creates an instance of the policy. With this feature, an Internet service provider (ISP) can create and revise a single policy that can apply to millions of users.

The BIG-IP system can enforce multiple levels of bandwidth limits through the dynamic policy. For example, a user could be limited by the maximum rate, a per user rate, and a per category rate (such as for an application), all from the same dynamic policy. When the total of the maximum user rate for all the instances exceeds the maximum rate specified in the dynamic policy, the BIG-IP system maintains fairness among all users and spreads the limitation equally among users belonging to a dynamic policy. In addition, you can specify per instance the maximum number of packets per second, which functions as a DoS (Denial of Service) limiter without fair share allocation.

You can also configure a dynamic bandwidth control policy to mark packets that exceed the maximum per-user rate for a specified session. The WAN router typically handles the marked packets. When marking is enabled, enforcement is implicitly disabled. You configure marking by using the **IP Marking (TOS/DSCP)** or **L2 Marking (802.1p)** setting. For example, a common use of QoS marking is for Voice over IP (VoIP) traffic. VoIP is usually assigned to the Expedited Forwarding (EF) class by using the DSCP value of 46, thus prioritized according to importance and sensitivity to loss/latency. You can mark packets per policy or per category (within a policy). Category marking supersedes policy marking.

The bandwidth controller is only an enforcer. For a dynamic bandwidth control policy, you also can use a virtual server (through iRules®), Policy Enforcement Manager™, or Access Policy Manager® to identify users and apply dynamic bandwidth control policies to traffic.

About the Maximum User Rate PPS setting

When you specify the value for the **Maximum User Rate PPS** setting for a dynamic bandwidth control policy, you are specifying the packets per second. Unlike the **Maximum Rate Per User**, this setting is not applied across a BIG-IP® system. This value depends on the packet sizes (MTU) configured in your network, and you need to tune the value accordingly. Although you can specify a lower value, the lowest recommended value is 256 KBPS. For a maximum user rate lower than 256 KPBS, the packet sizes in the network would need to be smaller than 1514 bytes. The issue is that for large packet sizes and a low maximum user rate, not enough bytes would be recharged for the bucket, and most packets would be dropped. If you are configuring categories, keep in mind that the maximum user rate is shared among the categories associated with a policy.

Task summary for creating a dynamic bandwidth control policy

Before you create a dynamic bandwidth control policy, F5 recommends that you select the **Source Address** for the **CMP Hash** setting on the VLAN properties screen for the VLAN that carries the traffic you want to manage. The BIG-IP® system uses source and destination hashes to control the way incoming traffic is distributed among the instances of the Traffic Management Microkernel (TMM) service. Subscriber-based

bandwidth control depends on having a unique one-to-one relationship between bandwidth control policy and subscriber. Subscribers are commonly identified using a unique IP address, and, therefore, load distribution among the instances of TMM service must use the source IP address as the key.

This screen snippet highlights the proper setting.

The screenshot shows a configuration window with a 'Configuration:' dropdown set to 'Advanced'. Below this is a table of settings:

Source Check	<input type="checkbox"/>
MTU	1500
Fail-safe	<input type="checkbox"/>
Auto Last Hop	Default
CMP Hash	Source Address
DAG Tunnel	Outer
DAG Round Robin	<input type="checkbox"/>

The 'CMP Hash' row is circled in red, highlighting the 'Source Address' selection.

Figure 4: CMP Hash setting for dynamic bandwidth control

This procedure describes the steps for attaching a dynamic bandwidth control policy to a traffic flow, and then applying the policy to traffic, using a virtual server. For information about using Policy Enforcement Manager™ to implement the policy, refer to the F5 documentation for Policy Enforcement Manager.

Task list

Creating a dynamic bandwidth control policy

Adding categories to a dynamic bandwidth control policy

Creating an iRule for a dynamic bandwidth control policy

Adding a dynamic bandwidth control policy to a virtual server

Creating a dynamic bandwidth control policy

You can create a dynamic bandwidth control policy to shape the traffic to which you apply the policy. You can configure the policy to mark packets per policy or per category. You can also specify the maximum number of packets per second per instance. Adding categories to a bandwidth control policy is a separate task, which you perform after you have created and saved the policy.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type a name for the bandwidth control policy.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for all the instances created for this dynamic policy.
The number must be in the range from 1 Mbps to 320 Gbps.
5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure to indicate the most bandwidth that each user or session associated with the bandwidth control policy can use.
The number must be in the range from 1 Mbps to 2 Gbps.
7. In the **Maximum User Rate PPS** field, type a number and select the unit of measure to specify the limit in packets per second that traffic is allowed per instance.

This is an optional setting, which functions as a DoS (Denial of Service) limiter without fair share allocation. You must also specify the **Maximum Rate Per User**. The system applies whichever value is lower to the traffic flow rate. When both values are specified, both must pass for packets to go through.

8. Enable the **Measure** setting, if you want to measure bandwidth on all future instances of this bandwidth control policy.

The system measures bandwidth with the frequency you specify in the **Log Period** setting, and sends it to the log publisher you specify using the **Log Publisher** setting.

9. From the **IP Marking (TOS/DSCP)** list, select **Specify** and type a number between 0 and 63 to assign a Type of Service (ToS) level to packets that exceed the maximum per-user rate.

If you do not want to set a ToS level, maintain the default setting, **Pass Through**.

10. From the **L2 Marking (802.1p)** list, select **Specify** and type a number between 0 and 7 to assign a Quality of Service (QoS) level to packets that exceed the maximum per-user rate.

If you do not want to set a QoS level, maintain the default setting, **Pass Through**.

11. Click **Finished**.

For the dynamic bandwidth control policy to take effect, you must attach the policy to a traffic flow, and then apply the policy to traffic, using a virtual server (through iRules®), Policy Enforcement Manager™, or Access Policy Manager®.

Adding categories to a dynamic bandwidth control policy

Before you can add categories, you must create a bandwidth control policy.

After you create a bandwidth control policy, you can add up to 32 categories of traffic for the policy to control. All the categories share the bandwidth specified for the bandwidth control policy, in accordance with the rate specified for each category.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click the name of the bandwidth control policy to which you want to add categories.
3. In the Categories area, click **Add**.
4. In the **Category Name** field, type a descriptive name for the category.
5. In the **Max Category Rate** field, type a value to indicate the most bandwidth that this category of traffic can use, and select the unit of measure from the list, or select **%** and type a percentage from 1 to 100.
If you specify a rate, the number must be in the range from 500 Kbps to the rate specified for the **Maximum Rate Per User** setting. A percentage indicates that this category can use up to the specified percentage of the maximum per-user rate. These values are upper limits (not minimum or guaranteed), so the sum can exceed the value you specified for the **Maximum Rate Per User** setting.
6. From the **IP Marking (TOS/DSCP)** list, select **Specify** and type a number between 0 and 63 to assign a Type of Service (ToS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a ToS level, maintain the default setting, **Pass Through**.
7. From the **L2 Marking (802.1p)** list, select **Specify** and type a number between 0 and 7 to assign a Quality of Service (QoS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a QoS level, maintain the default setting, **Pass Through**.
8. Click **Finished**.

Creating an iRule for a dynamic bandwidth control policy

To implement a dynamic bandwidth control policy, you can use iRules® to attach the policy to a user.

Note: For complete and detailed information iRules syntax, see the F5 Networks DevCentral web site (<http://devcentral.f5.com>).

1. On the Main tab, click **Local Traffic > iRules**.
The iRule List screen opens, displaying any existing iRules.
2. Click **Create**.
The New iRule screen opens.
3. In the **Name** field, type a unique name for the iRule.
The full path name of the iRule cannot exceed 255 characters.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.
For example, to apply the dynamic bandwidth policy `dynamic_bwc_policy200` to a user session, type the following iRule, where `set mycookie` defines a user session. A *session* is a combination of client IP address and port.

```
when CLIENT_ACCEPTED {
    set mycookie [IP::remote_addr]:[TCP::remote_port]
    BWC::policy attach dynamic_bwc_policy200 $mycookie
}
```

5. Click **Finished**.
The new iRule appears in the list of iRules on the system.

You have now identified the user for a dynamic bandwidth control policy.

You must then apply the iRule to the virtual server that intercepts the traffic you want to manage.

Adding a dynamic bandwidth control policy to a virtual server

After you create a dynamic bandwidth control policy and attach it to a flow or flows using iRules[®], you must apply the policy to traffic by adding the iRule to a virtual server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Resources**.
4. In the iRules area, click **Manage**.
5. From the **Available** list, select the name of the iRule that you want to assign, and using the Move button, move the name to the **Enabled** list.
6. Click **Finished**.

The BIG-IP[®] system now manages bandwidth for the traffic intercepted by this virtual server, according to the dynamic bandwidth policy specified in the assigned iRule.

Enforcing Bandwidth Control Provisioned by PCRF

Overview: Enforcing bandwidth control provisioned on PCRF

Policy Enforcement Manager™ (PEM) can enforce bandwidth limits provisioned by the PCRF using dynamic PCC rules. You do this by creating a dynamic bandwidth controller with the name `dynamic_spm_bwc_policy`. This bandwidth controller must be created on the BIG-IP system® using this predefined name. It does not need to be associated with an enforcement policy in PEM. Subscribers are assigned bandwidth in proportion to their configured rates. So for example, if subscriber A is assigned 4Mbps, and B is assigned 8Mbps, B will always get twice the bandwidth that A gets.

Task Summary

Creating a dynamic bandwidth control policy for PCRF

Creating a listener for subscriber discovery and policy provisioning

Creating a dynamic bandwidth control policy for PCRF

To set up bandwidth control through PCRF, you must have bandwidth control rules configured on the PCRF.

You can create a dynamic bandwidth controller so that PEM™ can enforce the maximum bit rate configured on the PCRF. You must follow the steps exactly as described here using the specified name for the bandwidth controller, and you must create the associated categories.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type the name `dynamic_spm_bwc_policy`.
4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for the resource you are managing.
The number must be in the range from 1 Mbps to 320 Gbps. This value is the amount of bandwidth available to all the connections going through this static policy.
5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure.
For example, use 50Mbps.
The number must be in the range from 1Mbps to 2Gbps. However, the value you use is just a place holder and is never used by the system. For this example, the value is overridden by the PCRF.
7. Enable the **Measure** setting, if you want to measure bandwidth on all future instances of this bandwidth control policy.
The system measures bandwidth with the frequency you specify in the **Log Period** setting, and sends it to the log publisher you specify using the **Log Publisher** setting.
8. Leave the **IP Marking (TOS/DSCP)** and **L2 Marking (802.1p)** values set to **Pass Through**, the default value.
9. Click **Finished**.

***Note:** After you finish configuring the bandwidth controller, the **Bandwidth Controllers** screen opens.*

10. Click the Bandwidth Control policy name, that you configured.
The Bandwidth Controllers policy page opens.
11. In the **Categories** field, add up to 32 categories of traffic that this bandwidth control policy manages.
All the categories share the specified bandwidth, in accordance with the rate specified for each category.

***Note:** Use the **Categories** setting only if you have not set values for the **IP Marking (TOS/DSCP)** or the **L2 Marking (802.1p)** setting.*

12. In the **Category Name** field, type a descriptive name for the category.
13. In the **Max Category Rate** field, type a value to indicate the most bandwidth that this category of traffic can use, and select the unit of measure from the list, or select **%** and type a percentage from 1 to 100.
If you specify a rate, the number must be in the range from 500 Kbps to the rate specified for the **Maximum Rate Per User** setting. A percentage indicates that this category can use up to the specified percentage of the maximum per-user rate. These values are upper limits (not minimum or guaranteed), so the sum can exceed the value you specified for the **Maximum Rate Per User** setting.
14. From the **IP Marking (TOS/DSCP)** list, select **Specify** and type a number between 0 and 63 to assign a Type of Service (ToS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a ToS level, maintain the default setting, **Pass Through**.
15. From the **L2 Marking (802.1p)** list, select **Specify** and type a number between 0 and 7 to assign a Quality of Service (QoS) level to packets that exceed the **Max Category Rate**.
If you do not want to set a QoS level, maintain the default setting, **Pass Through**.
16. Click **Finished**.

If this is the first bandwidth control policy created on a BIG-IP® device, the system also creates a default static bandwidth control policy named `default-bwc-policy` to handle any traffic that is not included in the policy you created. If you delete all bandwidth controllers, this policy is also deleted.

For PEM to enforce bandwidth control, you need to create a listener (Policy Provisioning and Online Charging Virtuals) with a Gx interface configured.

Creating a listener for subscriber discovery and policy provisioning

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. In the Policy Provisioning and Online Charging Virtuals area, click **Add**.
The New Configure Diameter Endpoint Provisioning and Online Charging screen opens.
3. In the **Name** field, type a unique name for the listener.
4. In the **Description** field, type a description of the listener.
5. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
6. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** or **Sd** from the **Supported Apps** options.

7. In the **Product Name** field, type the product name which is used to communicate with the PCRF.
8. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
9. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
10. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
11. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
12. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
13. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
14. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
15. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
16. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Implementation result

When traffic flows through the BIG-IP® system, the system limits the aggregated bandwidth for all subscribers to the **Maximum Rate** specified in the `dynamic_spm_bwc_policy` bandwidth control policy. The PCRF provides the **Maximum Rate Per User** for each subscriber, overriding the value in the bandwidth control policy. Policy Enforcement Manager™ restricts subscribers to the maximum user rate, and bandwidth is spread among subscribers fairly, on a best effort basis.

Configuring Tiered Services with Bandwidth Control

Overview: Configuring tiered services with bandwidth control

You can set up Policy Enforcement Manager to enforce different levels of bandwidth control on subscribers, providing more bandwidth to subscribers with higher tier subscriptions. Bandwidth control in this case is per subscriber and per application.

This implementation provides three tiers of service: gold (the highest level), silver (the next highest), and bronze (the lowest level). You create three dynamic bandwidth controllers, one for each tier to provide different bandwidth limits for subscribers with different plans. Each tier includes bandwidth control limits for three types of application traffic (P2P, audio-video, and web). You also create three enforcement policies, one for each tier. In the enforcement policies, rules applies bandwidth control to the different types of traffic.

Finally, subscribers are provisioned dynamically through a policy charging and rules function (PCRF) over a Gx interface. On the PCRF, you need to have associated subscribers with one of the subscriber tiers called gold, silver, and bronze.

Task Summary

Creating dynamic bandwidth control policies for tiered services

Creating enforcement policies for three tiers

Creating the rules for tiered bandwidth control

Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF

Creating dynamic bandwidth control policies for tiered services

You can create dynamic bandwidth controllers for tiered services so that PEM can enforce different rates of bandwidth control for subscribers having different policy levels. Use this procedure and the values specified to create three bandwidth controllers, one for each tier of service.

1. On the Main tab, click **Acceleration > Bandwidth Controllers**.
2. Click **Create**.
3. In the **Name** field, type the name of the bandwidth controller. In this example, name the three bandwidth controllers as follows:
 - Type `gold-bwc` for the premium subscription level.
 - Type `silver-bwc` for the medium subscription level.
 - Type `bronze-bwc` for the lowest subscription level.

There is no requirement to use these names, but it is convenient to use a similar name for the bandwidth controller and the enforcement policy that you will attach it to. Later in this example, you will attach the gold-bwc bandwidth controller to the gold enforcement policy.

4. In the **Maximum Rate** field, type a number and select the unit of measure to indicate the total throughput allowed for all the subscribers using each bandwidth controller. For this example, specify **10 Mbps** for all three bandwidth controllers

If you want to use different values, the number must be in the range from 1Mbps to 320Gbps.

5. From the **Dynamic** list, select **Enabled**.
The screen displays additional settings.
6. In the **Maximum Rate Per User** field, type a number and select the unit of measure relative to the tier of service. For example, use the following values:
 - For gold-bwc, specify **8 Mbps**.
 - For silver-bwc, specify **4 Mbps**.
 - For bronze-bwc, specify **2 Mbps**.
7. Leave the **IP Type of Service** and **Link Quality of Service** values set to **Pass Through**, the default value.
8. In the **Categories** field for each bandwidth controller, add three categories of traffic: P2P, Web, and Audio-video.
 - For gold-bwc, set P2P to 20%, set Web to 70%, and set Audio-video to 40%.
 - For silver-bwc, set P2P to 20%, set Web to 60%, and set Audio-video to 30%.
 - For bronze-bwc, set P2P to 20%, set Web to 50%, and set Audio-video to 20%.

In the rule for the policy, different types of traffic, P2P, web, or audio-video traffic is limited to a smaller percentage of the total bandwidth for all subscribers who use that policy.
9. Click **Finished**.

If this is the first bandwidth control policy created on a BIG-IP® device, the system also creates a default static bandwidth control policy named `default-bwc-policy` in the Common partition to handle any traffic that is not included in the policy you created. If you delete all bandwidth controllers, this policy is also deleted.

For PEM to enforce bandwidth control, you need to create enforcement policies with rules that refer to the bandwidth controller.

Creating enforcement policies for three tiers

To impose bandwidth control on multiple subscriber tiers, you need to create an enforcement policy for each tier. Use these steps and the values specified to create three enforcement policies.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.
 - Type `gold` for the premium subscription level.
 - Type `silver` for the medium subscription level.
 - Type `bronze` for the lowest subscription level.
4. Click **Finished**.
The new enforcement policy is added to the policy list.

You have three enforcement policies that represent the three tiers of subscriber traffic that you are creating.

After creating the enforcement policies, you add rules that specify how to treat the subscriber traffic in each tier. In the implementation being developed, subscribers in the different tiers will get different maximum amounts of bandwidth. Further limits will be placed on specific types of traffic (P2P, audio-video, and web).

Creating the rules for tiered bandwidth control

You next add rules to each of the enforcement policies you created (gold, silver, and bronze). The rules limit the amount of bandwidth that P2P, web, and audio-video traffic can use. Create three rules for each enforcement policy.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule. For the first rule, use the name **P2P**.
5. In the **Precedence** field, type an integer that indicates the precedence, 1 being the highest.
In this case, you can use any value, for example, 10, as the precedence for all the rules in all the policies because there is no conflict between the rules you are creating. Each rule applies to a different type of traffic: web, audio-video, and P2P.
6. In the **Classification** setting, specify the type of traffic.
 - a) For the first rule, from the **Category** list, select **P2P**. Use the default values for **Match Criteria (Match)** and **Application (Any)**.
 - b) Click **Add**.
7. In the **Rate Control** setting, for **Bandwidth Controller**, select the name of the bandwidth controller and category. Choose
 - a) For **Bandwidth Controller**, select the name that matches the policy you are working on. For example, if editing the gold policy, select **gold-bwc**.
 - b) For **Category**, select the category that matches the type of traffic specified by the name of the rule. For example, select **P2P**.
8. Click **Finished**.
9. Repeat steps 3-8 to create a second rule for audio-video traffic with these settings.

Option	Values
Name	Audio-video
Precedence	10
Classification Category	Audio_video
Rate Control-Bandwidth Controller	Same as the name of the policy you are adding the rule to (gold-bwc, silver-bwc, or bronze-bwc)
Bandwidth Controller-Category	Audio-video

10. Repeat steps 3-8 to create a third rule for web traffic with these settings.

Option	Values
Name	Web
Precedence	10
Classification Category	Web
Rate Control-Bandwidth Controller	Same as the name of the policy you are adding the rule to (gold-bwc, silver-bwc, or bronze-bwc)

Option	Values
Bandwidth Controller-Category	Web

The gold, silver, and bronze enforcement policies each have three rules called P2P, Web, and Audio-video. Each of the rules in the gold policy connects to the gold-bwc bandwidth controller; rules in the silver policy connect to the silver-bwc bandwidth controller and; rules in the bronze policy connect to the bronze-bwc policy.

Creating a listener for subscriber discovery with RADIUS and policy provisioning with PCRF

You create a listener to specify how to handle traffic for policy enforcement. Creating a listener does preliminary setup tasks on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting. You can also connect with a Policy and Charging Rules Function (PCRF) over a Gx interface.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. In the Policy Provisioning and Online Charging Virtuals area, click **Add**.
The New Configure Diameter Endpoint Provisioning and Online Charging screen opens.
3. In the **Name** field, type a unique name for the listener.
4. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

***Note:** When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.*

***Tip:** You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.*

The system will create a virtual server using the address or network you specify.

5. For the **Service Port** setting, type or select the service port for the virtual server.
6. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).
The system will create a virtual server for each protocol specified.
7. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. For subscriber provisioning using RADIUS, ensure that **Subscriber Identity Collection** is set to **RADIUS**.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. For the tiered services example, do not assign global policies.
12. To connect to a PCRF, from the **Diameter Endpoint** list, select **Enabled** and select **Gx** or **Sd** from the **Supported Apps** options.

13. In the **Origin Host** field, type the fully qualified domain name of the PCRF or external policy server, for example, `ocs.xnet.com`.
14. In the **Origin Realm** field, type the realm name or network in which the PCRF resides, for example, `xnet.com`.
15. In the **Destination Host** field, type the destination host name of the PCRF or external policy server, for example, `pcrfdest.net.com`.
16. In the **Destination Realm** field, type the realm name or network of the PCRF, for example, `net.com`.
17. For the **Pool Member Configuration** setting, add the PCRF servers that are to be members of the Gx endpoint pool. Type the **Member IP Address** and **Port** number, then click **Add**.
18. In the **Message Retransmit Delay** field, type the number of milliseconds to wait before retransmitting unanswered messages in case of failure from the BIG-IP system to the PCRF over the Gx interface. The default value is 1500.
19. In the **Message Max Retransmit** field, type the maximum number of times that messages can be retransmitted from the BIG-IP system to the PCRF. The default value is 2.
20. In the **Fatal Grace Time** field, type the time period in seconds that a diameter connection can remain disconnected before the system terminates all sessions associated with that diameter endpoint. The default value is 500.
21. Click **Finished**.
The Policy Enforcement Manager creates a listener.

When you create a listener, the Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. The system also creates a virtual server for the Gx interface with a diameter endpoint profile. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP system, the system handles policy enforcement.

Implementation result

When traffic flows through a BIG-IP® system, the system limits the aggregated bandwidth for all subscribers having a gold, silver, or bronze policy. Subscribers with a gold policy can use more of the total bandwidth than silver or bronze subscribers. Further, subscriber traffic in any of the tiers that is classified as audio-video, web, or P2P is limited to a percentage of the total bandwidth allowed for that tier.

For example, if a subscriber has a silver subscription level and PEM classifies their traffic as Web, the traffic is limited to 60% of the **Maximum Rate** specified in the `silver-bwc` bandwidth controller (4 Mbps). This leaves 2.4 Mbps as the maximum bandwidth for all web traffic of silver tier subscribers.

Configuring Service Chains

Overview: Configuring service chains

You can use the Policy Enforcement Manager™ to create service chains to route traffic to one or more value-added services on the way to its final destination. The *service chains* define the path and order that you want traffic to take. There are several value-added services involved and after each endpoint the traffic comes back to the BIG-IP system. An *endpoint* specifies each place you want to send the traffic, so the service chain is essentially between the value-added services endpoints for traffic to stop at on its way to the server it is headed to. For example, you can forward traffic sequentially for virus scanning, parental control, and caching.

You set up service chains by creating an enforcement policy that defines the traffic that you want to route to the service chain. Rules in the enforcement policy specify conditions that the traffic must match, and actions for what to do with that traffic. One of the actions you can take is to send the traffic to a service chain.

While a static service chain defines fixed value-added services, a dynamic service chain provides service chain action that can dynamically change depending on the flow of parameters and you can attach a steering policy that can override the decision of the next session. You can use dynamic service chain to insert or name header and steer different service. Internet Content Adaptation Protocol (ICAP) is one of the services possible to use in a service chain. Dynamic service chain makes the service chain intelligent and flexible by providing the following support:

- Ability to add or skip different value-added services endpoints by selecting policy based forwarding endpoint.
- Perform header insertion or removal per value-added service chain, depending on the policy.
- Includes one sideband value-added service in the service chain using ICAP as the protocol.

You can create listeners to set up virtual servers and associate enforcement policies with the traffic that is sent to them. The system also creates a Policy Enforcement profile that specifies the enforcement policy that the system uses for the service chain.

Task Summary

Creating a ICAP profile for policy enforcement
Creating a Request Adapt profile
Creating a Response Adapt profile
Creating an internal virtual server for ICAP server
Creating a pool
Creating endpoints for service chains
Creating dynamic service chains
Creating an enforcement policy
Configuring steering action policy
Adding rules to an enforcement policy
Creating a rule for forwarding traffic
Creating a data plane virtual group

About services profiles

You can configure the Internet Content Adaptation Protocol (ICAP) profile, request adaptation profile, and response adaptation profile for using the dynamic service chain feature in Policy Enforcement Manager™.

The internal virtual server references the pool of content adaptation servers. The internal virtual server also references an ICAP profile, which includes specific instructions for how the BIG-IP® system should modify each request or response. Once the request and response adapt profiles have been created, you can attach the profiles to the HTTP virtual server. The adapt profiles use multiple internal virtual servers for various content types.

The HTTP listener must have adapt profile set. The adapt profiles need to be configured as disabled and are enabled by PEM based on the policy action applied.

About service chain processing

The service chain endpoints that have steering policy attached, define the service chain. The dynamic service chain follows these processing strategies:

- The initial subscriber flow start processing of the service chain starts from the first service.
- The steering policy is evaluated before taking in account a default ICAP adaptation or the steering endpoint.

The steering policy changes the service chain in many ways:

- Skips the part of the service chain.
- Skips to different service of the ICAP or steering policy.
- Skip the rest of the service chain and route traffic to the network.
- Applies different services that are not on the chain. The steering policy can apply ICAP and skip the rest of the chain. It can also apply steering, skipping all ICAP on the VLAN. The service chain continues when the flow returns from the service.

Creating a ICAP profile for policy enforcement

You create this ICAP profile when you want to use an ICAP server to wrap an HTTP request in an ICAP message before the BIG-IP® system sends the request to a pool of web servers. The profile specifies the HTTP request-header values that the ICAP server uses for the ICAP message.

1. On the Main tab, click **Local Traffic > Profiles > Services > ICAP**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. Click **Finished**.

After you create the ICAP profile, you can assign it to an internal virtual server so that the HTTP request that the BIG-IP system sends to an ICAP server is wrapped in an ICAP message, according to the settings you specified in the ICAP profile.

Note: *Different services may require different ICAP profiles.*

Creating a Request Adapt profile

You create a Request Adapt type of profile when you want a standard HTTP virtual server to forward HTTP requests to an internal virtual server that references a pool of ICAP servers. A Request Adapt type of profile instructs the HTTP virtual server to send an HTTP request to a named internal virtual server for possible request modification.

1. On the Main tab, click **Local Traffic > Profiles > Services > Request Adapt**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. For the **Parent Profile** setting, retain the default value, `requestadapt`.
5. On the right-side of the screen, select the **Custom** check box.
6. Disable the setting by clearing the **Enabled** check box.
When you clear the **Enabled** check box, Policy Enforcement Manager™ controls this based on the policy.
7. In the **Preview Size** field, type a numeric value.
This specifies the maximum size of the preview buffer. This buffer holds a copy of the HTTP request header and the data sent to the internal virtual server, in case the adaptation server reports that no adaptation is needed. Setting the preview size to 0 disables buffering of the request and should only be done if the adaptation server always returns a modified HTTP request or the original HTTP request.
8. For the **Allow HTTP 1.0** setting, select the **Enabled** check box.
9. Click **Finished**.

After you perform this task, the BIG-IP® system contains a Request Adapt profile that a standard HTTP virtual server can use to forward an HTTP request to an internal virtual server for ICAP traffic. You need to attach a Request Adapt profile to a standard HTTP virtual server to forward the HTTP requests.

Creating a Response Adapt profile

You create a Response Adapt type of profile when you want a standard HTTP virtual server to forward HTTP responses to an internal virtual server that references a pool of ICAP servers. A Response Adapt type of profile instructs the HTTP virtual server to send an HTTP response to a named internal virtual server for possible response modification.

1. On the Main tab, click **Local Traffic > Profiles > Services > Response Adapt**.
2. Click **Create**.
3. In the **Name** field, type a unique name for the profile.
4. For the **Parent Profile** setting, retain the default value, `responseadapt`.
5. On the right-side of the screen, select the **Custom** check box.
6. Disable the setting by clearing the **Enabled** check box.
When you clear the **Enabled** check box, Policy Enforcement Manager™ controls the profile based on the policy.
7. In the **Preview Size** field, type a numeric value.
This specifies the maximum size of the preview buffer. This buffer holds a copy of the HTTP response header and the data sent to the internal virtual server, in case the adaptation server reports that no adaptation is needed. Setting the preview size to 0 disables buffering of the response and should only be done if the adaptation server always returns a modified HTTP response or the original HTTP response.
8. For the **Allow HTTP 1.0** setting, check the **Enabled** check box.

After you perform this task, the BIG-IP® system contains a Response Adapt profile that a standard HTTP virtual server can use to forward an HTTP response to an internal virtual server for ICAP traffic. You need to attach a Response Adapt profile to a standard HTTP virtual server to forward the HTTP responses.

Creating an internal virtual server for ICAP server

You perform this task to create a standard virtual server that can forward an HTTP request or response to an internal virtual server. The internal virtual server then sends the request or response to a pool of ICAP servers before the BIG-IP® system sends the request or response to the client or web server.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. From the **Type** list, select **Internal**.
5. For the **State** setting, verify that the value is set to **Enabled**.
6. From the **Configuration** list, select **Advanced**.
7. From the **ICAP Profile** list, select the name of the HTTP profile that you created previously.
8. From the **Source Address Translation** list, select **Auto Map**.
The BIG-IP® system uses all of the self IP addresses as the translation addresses for the pool.
9. Optionally, from the **OneConnect Profile** list, select a custom OneConnect profile.

Note: Setting **OneConnect Profile** to ICAP virtual server, is highly recommended when configuring ICAP virtual.

10. From the **Default Pool** list, select the pool of ICAP servers that you previously created.
11. Click **Finished**.

After you create the virtual server, the BIG-IP® system can forward an HTTP request or response to a pool of ICAP servers before sending the request or response to the client or web server, respectively.

Creating a pool

You can create a pool of servers that you can group together to receive and process traffic.

1. On the Main tab, click **Local Traffic > Pools**.
The Pool List screen opens.
2. Click **Create**.
The New Pool screen opens.
3. In the **Name** field, type a unique name for the pool.
4. Using the **New Members** setting, add each resource that you want to include in the pool:
 - a) (Optional) In the **Node Name** field, type a name for the node portion of the pool member.
 - b) In the **Address** field, type an IP address.
 - c) In the **Service Port** field, type a port number, or select a service name from the list.
 - d) (Optional) In the **Priority** field, type a priority number.
 - e) Click **Add**.
5. Click **Finished**.
6. Repeat these steps for each pool you want to create.

The new pool appears in the Pools list.

Creating endpoints for service chains

Before you can create an endpoint, you need to create a pool that specifies where you want to direct the classified traffic.

If you plan to set up a service chain, you need to create one or more endpoints that specify the locations of the value-added services to which to send the traffic.

1. On the Main tab, click **Policy Enforcement > Forwarding > Endpoints**.
The Endpoints screen opens.
2. Click **Create**.
The New Endpoint screen opens.
3. In the **Name** field, type a name for the endpoint.
4. From the **Pool** list, select the pool to which you want to steer a particular type of traffic.
5. Use the default values for the other fields.
6. Click **Finished**.
The endpoint you created is on the endpoint list.

You link the endpoints together by creating a service chain.

Creating dynamic service chains

Before you can create a service chain, you need to have created endpoints for every service that you want the traffic to be directed to. Set up the servers at those endpoints to handle the traffic and (if conditions are right), return it to the BIG-IP® system. You should have attached the HTTP virtual server to the request adapt profile and response adapt profile. You also need to create VLANs for every traffic entry point.

To send traffic to multiple endpoints, including value-added services, you create service chains that define where to send traffic on the way to its final destination. This way, the system can route traffic to other servers that can handle additional functions. Additionally, you can attach a steering action policy, such as modify headers, when you create a service chain which can be later modified at the other end.

Note: *If you want to use steering policy, you must define endpoint in service chain.*

1. On the Main tab, click **Policy Enforcement > Forwarding > Service Chains**.
The Service Chains screen opens.
2. Click **Create**.
The New Service Chains screen opens.
3. In the **Name** field, type a name for the service chain.
4. In the **Service Chain List** setting, add the endpoints to the service chain. For each place you want to send the traffic, specify the following information:
 - a) From the **Service Endpoint Name** list, type the name of the service endpoint where the traffic is going to.
 - b) From the **VLAN** list, select the name of the VLAN where the traffic is coming from.

Note: *Your first service chain should have subscriber VLAN in the VLAN field.*

- c) From the **Policy** list, select the name of the steering policy.

Note: If all the service endpoints do not have a steering policy, the service chain is static.

Important: If the policy defining the steering does not match the policy set in the service chain, then the service chain is not processed.

- d) From the **Forwarding Endpoint** list, select the name of the endpoint to which you send traffic.

When you configure a new forwarding endpoint (**Policy Enforcement > Forwarding > Endpoints**), set **Address Translation** and **Port Translation** as **Disabled**.

Note: You need to always configure a default forwarding endpoint or else the flow will exit the service chain and get skipped. If you are in the final leg, then configure without default.

Important: When you use ICAP service, you cannot have a ICAP and a forwarding endpoint on the same service endpoint.

- e) From the **Service Option** list, select the service option in case the service endpoint is not reachable. Select **Optional** if you want to skip the service endpoint. Select **Mandatory** if you want all traffic flows dropped.
-

Note: To use dynamic service chain, select **Optional**. If service endpoint is not available and set to mandatory, you cannot steer policies.

Note: The **Service Option** parameter works only if the right endpoint has a monitor set in the pool. For example, set gateway ICMP to the pool. Otherwise, traffic is dropped even if **Optional** is set.

- f) From the **Internal Virtual** list, select the internal ICAP virtual server.
-

Important: You cannot have consecutive ICAP on the same VLAN.

- g) Click **Add**.

5. From the **ICAP Type** list, select the action you want to implement. Select **Request** to send only HTTP requests to ICAP server. Select **Response** to send only HTTP responses to ICAP server. Select **Request and Response** to have both requests and responses.

- Select **Response** to send only HTTP responses to ICAP server.
 - Select **Request and Response** to have both requests and responses.
-

Note: Select the **Internal Virtual** to configure the **ICAP Type** setting.

6. Click **Finished**.
-

Note: If steering action is applied after the ICAP request, service endpoint with forwarding endpoint should have the same VLAN configured as the service endpoint with ICAP enabled.

You can direct traffic to the service chain you created in the policy rules in an enforcement policy.

Creating an enforcement policy

If you want to classify and intelligently steer traffic, you need to create an enforcement policy. The policy describes what to do with specific traffic, and how to treat the traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click **Create**.
The New Policy screen opens.
3. In the **Name** field, type a name for the policy.

***Tip:** When creating policies you plan to apply globally or to unknown subscribers, it is a good idea to include the word `global` or `unknown` in the policy name to distinguish these from other subscriber policies.*

4. From the Transactional list, select **Enabled** if you want the BIG-IP system to allow policy enforcement on each HTTP transaction.
5. Click **Finished**.

***Important:** The system performance is significantly affected, depending on complexity of the classification and the type of policy action.*

The new enforcement policy is added to the policy list.

Now you must add rules to the enforcement policy to define traffic filters and actions.

Configuring steering action policy

You can configure HTTP headers of the steering policy in the BIG-IP® system.

***Note:** If the steering action is enabled, steering policy is evaluated based on the VLAN flow. If no steering policy is configured, then the default endpoint is the next service endpoint.*

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Tip:** All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.*

6. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.
More modify header configuration options display.
7. To modify the HTTP request header, select the action you want to implement.
 - Select **Insert String Value** to insert a stringvalue that you have specified before.

- Select **Insert Value from Script** to specify that the BIG-IP system can insert value received from the TCL expression.
 - Select **Remove** to remove the string value that you previously created.
8. In the **Header Name** field, type a header name.
 9. In the **String Value** field, type a string value for the header.
 10. Click **Finished**.

You can add more rules to an enforcement policy in addition to configuring HTTP header action.

Adding rules to an enforcement policy

Before you can add rules to an enforcement policy, you need to create the policy, then reopen it.

You add rules to an enforcement policy to select the traffic you want to affect, and the actions to take. A *rule* associates an action with a specific type of traffic. So you can, for example, add a rule to select all audio-video traffic and send it to a pool of servers that are optimized to handle that type of traffic.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. From the **Modify Header** list, select **Enabled**, to modify the HTTP request header.
More modify header configuration options display.
8. Use the Reporting, Quota, Forwarding, Modify Header or QoS areas to specify what you want to do with the traffic that you are classifying or specify what actions you want to apply to the traffic.
Other tasks describe how to do this in detail.
If you leave **Gate Status** enabled (default) and specify no other actions, the system stores traffic classification statistics on the BIG-IP system, and forwards the traffic to its destination without any further action.
9. From the **Congestion Detection** list, select **Enable**, to congestion detection in the Radio Access Network.
 - a) In the **Threshold** field, type the lower threshold bandwidth for a session. The default value is 1000kbs.
 - b) For **Destination** list, select the publisher name from the HSL publisher drop-down list.

The state of congestion detection is now controlled by policy application, and different subsets of subscribers can have different settings. This enables congestion-detection for specific types of applications as it pairs with specific policy rule conditions.

10. Click **Finished**.

11. Repeat steps 3-8 to create as many rules as needed to handle the traffic you are interested in.

The enforcement policy includes the rules with the conditions and actions you added.

Now you need to associate the enforcement policy with the virtual server (or servers) to which traffic is directed.

Creating a rule for forwarding traffic

You can create a rule that forwards traffic to an endpoint. For example, you might want to direct video traffic to a server that is optimized for video viewing.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. Use the Classification, URL, Flow, and Custom Criteria tabs to identify the traffic that you want to be affected by this rule.
7. In the Gate area, for **Gate Status**, select **Enabled**.
Options provide several ways to forward the traffic.
8. In the Forwarding area, for **HTTP Redirect**, select **Enabled**, and type the URL.
9. From the Forwarding list, select an option where you would like to forward the traffic.

Options	Description
Route to Network	The traffic flow is forwarded to the default destination.
Forwarding to Endpoint	The flow is steered to a different destination and you can select one of the endpoints.
Forward to ICAP virtual Server	The flow is forwarded to the ICAP virtual server.

10. From the **Forwarding Fallback Action** list, select **Drop** or **Continue** to specify if the connection can remain unchanged or should be dropped if the forwarding action fails.

11. From the **ICAP Virtual Server** list, select an internal virtual server that you have created, or click **Create** to create a new internal virtual server.
12. From the **ICAP Type** list, select an ICAP adaptation type.
 - Select **Request** to send a portion of the request to the ICAP server.
 - Select **Response** to receive a portion of the response from the ICAP server.
 - Select **Request** and **Response** to have both types of adaptation.
13. From the **Service Chain** list, select **Create** to direct traffic to more than one location (such as value-added services).
14. Click **Finished**.

You have created a rule that forwards traffic.

Creating a data plane virtual group

If you want to steer specific traffic (or otherwise regulate certain types of traffic) you must first develop appropriate enforcement policies. If using a Gx interface to a PCRF, you need to create a new virtual group in listeners that connect to a PCRF.

You can create listeners that specify how to handle traffic for policy enforcement. Creating a listener performs preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, and reporting.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Data Plane**.
The Data Plane screen opens.
3. Click **Add Group**.
The New Virtual Group screen opens.
4. In the **Name** field, type a unique name for the listener.
5. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

6. For the **Service Port** setting, type or select the service port for the virtual server.
7. From the **VLAN and Tunnel Traffic** list, select **Enabled on**. Then, for the **VLANs and Tunnels** setting, move the VLAN or VLANs on which you want to allow the virtual servers to share traffic from the **Available** list to the **Selected** list.
8. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
9. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

Note: For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

10. Click Finished.

The Policy Enforcement Manager creates a listener.

When you create a listener, Policy Enforcement Manager™ also creates virtual servers for each type of traffic (TCP, UDP, or both and IP), and a virtual server for HTTP traffic. The system sets up classification and assigns the appropriate policy enforcement profile to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Now you can send traffic through the network. As network traffic moves through the BIG-IP® system, the system classifies the traffic, and if you have developed policies, the system performs the actions specified by the enforcement policy rules.

Provisioning Dynamic Subscribers

Overview: Provisioning dynamic subscribers

If you have subscribers that are managed on a separate policy charging and rules function (PCRF), you can connect the BIG-IP® system to that policy server to provision dynamic subscribers. *Dynamic subscribers* are subscribers that are managed by a PCRF.

The BIG-IP system receives traffic from GGSN, a gateway between the GPRS mobile network and the Internet. When a subscriber makes a request that is routed to the BIG-IP system, the Policy Enforcement Manager™ queries the PCRF over a Gx interface. The PCRF responds with information about the subscriber. This information is stored on the BIG-IP system, which recognizes the subscriber in future requests.

You can use dynamic subscriber provisioning alone or in combination with static subscribers.

Provisioning dynamic subscribers

If you want to steer specific traffic, or otherwise regulate certain types of traffic, you need to have developed enforcement policies.

To provision subscribers dynamically through a separate PCRF, you create listeners that specify how to handle traffic for policy enforcement. Creating a listener provides preliminary setup on the BIG-IP® system for application visibility, intelligent steering, bandwidth management, reporting, and other features.

1. On the Main tab, click **Policy Enforcement > Listeners**.
The Listeners screen opens.
2. Click **Data Plane**.
The Data Plane screen opens.
3. Click **Add Group**.
The New Virtual Group screen opens.
4. In the **Name** field, type a unique name for the listener.
5. In the **Destination Address** field, type the IP address of the virtual server. For example, 10.0.0.1 or 10.0.0.0/24.

Note: When you use an IPv4 address without specifying a prefix, the BIG-IP® system automatically uses a /32 prefix.

Tip: You can use a catch-all virtual server (0.0.0.0) to specify all traffic that is delivered to the BIG-IP® system. Configure the source and destination setting, during forwarding mode only. In the relay mode, the client does not have an IP address and the DHCP provides the client with an IP address.

The system will create a virtual server using the address or network you specify.

6. For the **Service Port** setting, type or select the service port for the virtual server.
7. From the **Protocol** list, select the protocol of the traffic for which to deploy enforcement policies (**TCP**, **UDP**, or **TCP and UDP**).

The system will create a virtual server for each protocol specified.

8. To use network address translation, from the **Source Address Translation** list, select **Auto Map**.
The system treats all of the self IP addresses as translation addresses.
9. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor from the **Available** list to the **Selected** list.
10. For the **VLANs and Tunnels** setting, move the VLANs and tunnels that you want to monitor for RADIUS traffic from the **Available** list to the **Selected** list.
11. In the Policy Provisioning area, select enforcement policies to apply to the traffic.
 - a) For **Global Policy**, move policies to apply to all subscribers to **High Precedence** or **Low Precedence**.

***Note:** For URL categorization to take effect, you need to associate the enforcement policy with a classification profile.*

- b) For **Unknown Subscriber Policy**, move policies to use if the subscriber is unknown to **Selected**.

The system applies the global policy to all subscribers in parallel with the subscriber policies, and must be configured with unknown subscriber policy. High-precedence global policies override conflicting subscriber policies, and low-precedence policies are overridden by conflicting subscriber policies.

When you create a listener, the Policy Enforcement Manager[®] also creates virtual servers for each type of traffic (TCP, UDP, or both), and a virtual server for the Gx interface. The system also creates a virtual server to handle HTTP traffic. The system assigns the appropriate classification and policy enforcement profiles to the virtual servers. If you are connecting to a RADIUS authentication server, a virtual server for RADIUS is also added.

Provisioning Static Subscribers

Overview: Provisioning static subscribers

If you have subscribers that are not managed on a separate policy charging and rules function (PCRF), you can create static subscribers. *Static subscribers* are individual subscribers of services that are not managed by a separate policy server. You can add static subscribers directly to the Policy Enforcement Manager™, and assign enforcement policies to them. Each subscriber can have one or more enforcement policies associated with them.

You can use static subscribers alone or in combination with those managed by a PCRF.

Provisioning multiple subscribers

Before you create a subscriber, you need to create the enforcement policy that you want to associate with the subscriber.

You can provision multiple subscribers to BIG-IP® Policy Enforcement Manager™ and associate policies with them. This is useful, for example, if you are not using a separate policy charging and rules function (PCRF) for subscribers, or if you need to add a subscriber with a unique enforcement policy.

1. On the Main tab, click **Policy Enforcement > Subscribers**.
The Subscribers screen opens.
2. Click **Create**.
The New Subscriber screen opens.
3. From the **Subscriber ID Type** list, select the format of the subscriber ID.

Option	Description
E.164	A numbering plan that defines the format of an MSISDN international phone number (up to 15 digits). The number typically consists of three fields: country code, national destination code, and subscriber number.
IMSI	International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.
NAI	Network Access Identifier. A fully qualified network name (FQDN) in the form <user>@<realm>; identifies a subscriber and the home network to which the subscriber belongs.
Private	The subscriber ID type is private for the given deployment.

4. In the **Subscriber ID** field, type a unique identifier (up to 64 characters) for the subscriber, such as a phone number. The subscriber ID type determines the format.
5. In the **Subscriber IP Address** area, type the IP address of the subscriber in IPv4 or IPv6-Prefix format. Click **Add** to add multiple IP addresses. This field is optional but recommended.

Tip: Assigning the subscriber IP addresses ensures that the subscriber gets the entitled service faster.

6. In the **Policies** setting, select at least one enforcement policy from the **Available** list and move it to the **Selected** list.

The selected policy is the one that the system enforces for the subscriber you are adding.

***Note:** You can assign a transactional policy to an active subscriber if you have created a transactional policy with the transactional reporting rule action.*

7. Click **Finished**.

Policy Enforcement Manager creates a static subscriber.

When the subscriber accesses the network through the BIG-IP system, Policy Enforcement Manager applies the policy you assigned to the subscriber traffic.

Subscriber CSV file format

You can upload a list of static subscribers for policy enforcement on the BIG-IP® system. A CSV file is a text file in a comma-separated-values format. Microsoft® Excel® can be used to view and create such files. The list of subscribers must be in a text file in comma separated value (CSV) format.

Each line of the CSV file represents one subscriber in the format:

```
subscriber_ID,subscriber_ID_type,number_of_IP_addresses,subscriber_IP,policy_1,policy_2,policy_n
```

subscriber ID

A unique identifier for the subscriber that depends on the subscriber ID type.

subscriber ID type

The format of the subscriber ID. This field is optional, but the comma is required. If you omit subscriber ID type, the system assigns the default value of IMSI (International Mobile Subscriber Identity).

subscriber IP

The IP address of the subscriber. This field is optional, but the comma is required.

subscriber IP list

The list of subscribers. This field is optional, but the comma is required.

policy_1,policy_2,policy_n

One policy or more policies assigned to the subscriber. Multiple policies must be separated by commas.

***Note:** To allow multiple IPs you need to change the DB variable (`tmm.pem.session.ip.addr.max`).*

For example:

```
subscriber_11.1.1.1,private,11.1.1.1,bronze,gold,silver
```

You must include Subscriber ID and at least one policy enforcement policy for each subscriber. You need to include the comma for missing fields. Do not include spaces between values. The policies listed must be included on the policies list in **Policy Enforcement > Policies**, and be provisioned using a listener or a policy enforcement profile.

For example, to specify subscriber 8315555555 in IMSI format with a gold policy, you can leave out the subscriber ID type and subscriber IP address:

```
8315555555,,gold
```

Provisioning a file of static subscribers

You can upload a CSV file containing a list of static subscribers. The file must be in a specific CSV file format.

1. On the Main tab, click **Policy Enforcement > Subscribers**.
The Subscribers screen opens.
2. Click **Upload**.
3. In the **File Name** setting, click the button to browse to the text file.
The system opens a screen where you can select the text file
4. Click **Upload**.

The Policy Enforcement Manager™ uploads the static subscriber file in chunks of 1000 subscribers. The system performs a validation on each chunk. If a validation fails, the subscribers in the current chunk and subsequent chunks are not imported. However, the subscribers loaded in previous chunks are imported onto the system. The Policy Enforcement Manager applies the appropriate policy to traffic from the subscribers in the list.

About Subscriber and Policy Provisioning

Overview: Subscriber and policy provisioning

The Policy Enforcement Manager™ (PEM™) is flexible in extracting Attribute Value Pair (AVPs) from Radius and encoding AVPs to and from PCRF. This solves interoperability issues to a larger extent. The configuration includes creation of RADIUS and Gx protocol profiles. The RADIUS AVPs, diameter AVPs, as well as manually defined parameters can be stored as subscriber attributes in the BIG-IP® system.

Subscriber and policy provisioning addresses the following factors:

- Different AVPs in RADIUS accounting messages are extracted and mapped to subscriber attributes for the session. This is used for subscriber identity, additional subscriber characterization related to network type, location and so on, and their changes. For example, RAT-Type, 3GPP-Location-Info, 3GPP-SGSN-Address, or allowing discovery of certain events included and covered by the accounting process.
- Add AVPs specified in CCR-I, CCR-u, CCR-T and RAA messages and also extract AVPs from CCA-I, CCA-U and RAR. For example, one can configure to send multiple Subscription-ID read over RADIUS to PCRF.
- Custom policy decisions from AVPs can be made with use of extracted parameters. For example, Operator can configure policy based on RAT-type (2G/3G/4G) or any custom AVP extracted of the subscriber.
- Extracted parameters can be inserted in to reporting records over HSL and Gx.

Task summary

Configuring RADIUS in PEM profile

Configuring RADIUS AVP in PEM profile

Configuring RADIUS message

Configuring Gx in PEM profile

Configuring Gx message

Configuring Subscriber Attributes in PEM profile

Configuring Diameter AVP in PEM profile

Configuring RADIUS in PEM profile

You can create a custom PEM RADIUS protocol profile and add a message defining how the RADIUS message is to be processed in the Policy Enforcement Manager. Radius messages configured on subscribers can be in, out or any direction.

1. On the Main tab, click **Policy Enforcement > Profiles > Protocol > RADIUS**.
The Configuration screen opens.
2. Click **Create**.
The New RADIUS Profile screen opens.
3. In the **Name** field, type a unique name for the protocol.
4. In the **Description** field, type optional descriptive text for the protocol.
5. From the Subscriber ID Type list, select an option which the identifier represents when the subscriber session is created.

Options	Subscriber ID Type
E164	A number that defines the format of an MSISDN international phone number (up to 15 digits).
NAI	A fully qualified network name that identifies a subscriber and the home network to which the subscriber belongs.
IMSI	A globally unique code number, that identifies a GSM, UMTS, or LTE mobile phone user.
Private	The subscriber ID type is private for the given deployment.

- From the **Subscriber ID List** list, in the **Subscriber ID Name** setting, type the subscriber ID name.
Enabled.
- From the **Subscriber ID List** list, in the **Order** setting, type the order of RADIUS AVPs when constructing the subscriber ID.
- From the **Subscriber ID List** list, in the **RADIUS AVP** setting, select the value of RADIUS AVP which you have used to create the subscriber ID.
- From the **Subscriber ID List** list, in the **Prefix** setting, type the suffix string to be added to THE extracted subscriber ID attribute that is specified in the RADIUS AVP for the subscriber session created.
- From the **Subscriber ID List** list, in the **Suffix** setting, type the suffix string when constructing subscriber ID with the value of the RADIUS AVP.
- To configure the RADIUS Message, click **Create & Add RADIUS Message**.
The Policy Enforcement Manager creates a new RADIUS message page, where you can configure the AVP List.
- Click **Finished**.

You have created a custom RADIUS profile.

Configuring RADIUS AVP in PEM profile

You can create a RADIUS AVPs that could be used to extract RADIUS attributes for subscriber sessions. PEM already has a list of predefined standard AVP attributes already defined as part of configuration.

- On the Main tab, click **Policy Enforcement > Profiles > Protocols > RADIUS AVP**.
The RADIUS AVP screen opens.
- Click **Create**.
The Subscriber Attribute Properties screen opens.
- In the **Name** field, type a unique name for the profile.
- In the **Description** field, type optional descriptive text for the profile.
- From the **Data Type** list, select an option for the data type of the RADIUS AVP.

Options	Data Type
3GPP RAT Type	The value format to be encoded or decoded as the 3GPP-RAT-Type defined in 3GPP TS 29.061.
3GPP User Location Information	The value format to be encoded or decoded as the 3GPP-User-Location-Info defined in 3GPP TS 29.061.
IP Address	The Account Status Type AVP is set to 3 (Interim-Update).
IPv4 Address	The IPv4 address in network byte order.
IPv6 Address	The IPv6 address in network byte order.

Options	Data Type
Integer	The 32-bit unsigned integer in network byte order.
Octet	The TF-8 text [RFC3629], totaling 253 octets or less in length.
String	The binary data, totaling 253 octets or less in length. This includes the opaque encapsulation of data structures defined outside of RADIUS.
Time	The 32-bit unsigned value in network byte order and in seconds since 00:00:00 UTC, January 1, 1970.

6. In the **Minimum Length** field, type the minimum data length of the RADIUS AVP. The default value is 0.
7. In the **Maximum Length** field, type the maximum data length of the RADIUS AVP. The default value is 0.
8. In the **Vendor ID Length** field, type the vendor ID of the RADIUS VSA. Type the default value is 1045.
9. In the **Vendor Type** field, type the vendor type of the RADIUS VSA. The default value is 20.
10. In the **Type** field, type of the RADIUS AVP. The default value is 26.
11. Click **Finished**.

You have created RADIUS AVP that help with policy decisions.

Configuring RADIUS message

You can configure the RADIUS message in the previously created Policy Enforcement Manager™ (PEM) RADIUS protocol profile, to extract AVPs for the subscriber session creation in the ingress direction. PEM provides the list of RADIUS messages which are populated with well known standard AVPs that are extracted to create attributes, in the subscriber session. To apply mapping between RADIUS AVPs TO PEM subscriber attributes, create mapping between each RADIUS AVP.

1. In the **Name** field, type a unique name for the protocol.
2. From the **Direction** list, in the **AVP List** setting, select **Any**, **In** or **Out** to process the radius message in both ingress and egress, ingress or egress direction respectively.
3. From the Message Type list, select an option which the identifier represents when the subscriber session is created.

Options	Message Type
Accounting Request Start	The Account Status Type AVP is set to 1 (Start).
Accounting Request Stop	The Account Status Type AVP is set to 2 (Stop).
Accounting Request Interim Update	The Account Status Type AVP is set to 3 (Interim-Update).

4. To apply mapping between RADIUS AVPs and PEM subscriber attributes configure the actions you want to implement.
 - In the **AVP** field, type the name of the application service to which the AVP belongs.
 - In the **Default** field, type the default value that is used in the subscriber session, if the RADIUS message is not present.
 - From the **Ingress** list, select the **Import** option for the RADIUS AVP to be parsed and the value to be stored in the subscriber attribute. The default value is **None**.

- From the **RADIUS AVP** list, select the name of the RADIUS AVP. The default value is **None**.
- From the **Subscriber Attribute** list, select the name of the subscriber session attribute to be mapped to RADIUS AVP. The default value is **None**.

5. Click **Finished**.

You have created a custom RADIUS message.

Configuring Gx in PEM profile

You can create a custom GX protocol profile and add a message which is defined as the Gx Credit-Control-Request (CCR), in any direction.

1. On the Main tab, click **Policy Enforcement > Profiles > Protocol > Gx**.
The Configuration screen opens.
2. Click **Create**.
The New Gx Profile screen opens.
3. In the **Name** field, type a unique name for the protocol.
4. In the **Description** field, type optional descriptive text for the profile.
5. From the Subscriber ID Type list, select an option which the identifier represents when the subscriber session is created.

Options	Subscriber ID Type
E164	A number that defines the format of an MSISDN international phone number (up to 15 digits).
NAI	A fully qualified network name that identifies a subscriber and the home network to which the subscriber belongs.
IMSI	A globally unique code number, that identifies a GSM, UMTS, or LTE mobile phone user.
Private	The subscriber ID type is private for the given deployment.

6. From the **Diameter AVP** list, in the **Subscriber ID** setting, select the diameter AVP.
7. From the **Type AVP** list, in the **Subscriber ID** setting, select the AVP that is specified in message that should be matched.
8. To configure the Gx Message, click **Create & Add Gx Message**.
The Policy Enforcement Manager creates a new Gx protocol profile message.
9. Click **Finished**.

You have created a custom Gx profile.

Configuring Gx message

You can configure the Gx message in the Policy Enforcement Manager. The message is defined as RADIUS accounting on the ingress direction.

1. In the **Name** field, type a unique name for the protocol.
2. From the **Direction** list, in the **AVP List** setting, select **Any**, **In** or **Out** to process the radius message in both ingress and egress, ingress or egress direction respectively.
3. From the Message Type list, select the message type.

4. From the **AVP List** setting, apply mapping of Diameter AVPs to subscriber session attribute for specific Gx message, by configuring the following:

- In the **AVP Name** field, type the name of the application service to which the AVP belongs.
- In the **Default** field, type the diameter AVP default value.
- From the **Protected Flag** list, select **Enabled** to choose the value of the protected flag, in the diameter AVP, when the message is inserted. This flag only applies to diameter AVP in outgoing message.

***Note:** The parent AVP inherits flags of child AVPs.*

- From the **Mandatory Flag** list, select **Enabled** to choose the value of the mandatory flag, in the diameter AVP, when the message is inserted. This flag only applies to diameter AVP in outgoing message.
- From the **Vendor-Specific Flag** list, select **Enabled** to choose the value of the vendor-specific flag, in the diameter AVP, when the message is inserted. This flag only applies to diameter AVP in outgoing message.
- From the **Diameter AVP** list, select the name of a configured diameter AVP. The default value is **None**.
- In the **Parent Label** field, type the name of a parent label which groups AVPs that can be combined.

***Note:** The AVPs with the same parent-label are combined in the same grouped AVP.*

- From the **Subscriber Attribute** list, select the name of a configured subscriber session attribute. The default value is **None**.
- From the **Include Interim Message** list, select **Enabled** for the AVP to be included in the interim-message (ccr-u only) updates which are generated if there is any change related to session parameters.

***Note:** This flag only applies to Diameter AVP in outgoing message.*

- From the **Include Reporting Message** list, select **Enabled** for the AVP to be included in the reporting message (ccr-u only) updates which are generated for reporting usage information.

***Note:** This flag only applies to Diameter AVP in outgoing message.*

5. Click **Finished**.

You have created a custom Gx message.

Configuring Subscriber Attributes in PEM profile

You can configure subscriber attributes in PEM profile. The subscriber attributes are used for storing values imported from RADIUS or DIAMETER and manually or iRule set of values for further use.

1. On the Main tab, click **Policy Enforcement > Profiles > Protocols > Subscriber Attributes**. The Subscriber Attributes screen opens.
2. Click **Create**. The Subscriber Attribute Properties screen opens.
3. In the **Name** field, type a unique name for the profile.
4. In the **Description** field, type optional descriptive text for the profile.
5. From the **Import** list, select **Enabled** for the subscriber attribute to be imported (parsed) from the incoming messages. The default value is **None**.

6. From the **Export** list, select **Enabled** for the subscriber attribute to be exported (inserted) to the outgoing messages. The default value is **Enabled**.
7. From the **Well Known Attribute ID** list, select an option for an identifier of a well-known (built-in) subscriber attribute.

***Note:** The system provides a special handling for well-known subscriber attributes. Session reporting records have the most well-known attributes by default.*

8. Click **Finished**.

You have created a subscriber attribute in PEM profile.

Configuring Diameter AVP in PEM profile

You can configure Diameter AVP, that can be used in Gx messages.

1. On the Main tab, click **Policy Enforcement > Profiles > Protocols > Diameter AVP**. The Diameter AVP screen opens.
2. Click **Create**. The Diameter AVP Properties screen opens.
3. In the **Name** field, type a unique name for the profile.
4. In the **Description** field, type optional descriptive text for the profile.
5. From the **Parent AVP** list, select **None** for name of the parent AVP, if it is in a grouped AVP.
6. In the **AVP Code** field, type the AVP code of the diameter AVP.
7. From the **Data Type** list, select the data type of the diameter AVP. The default value is **Octet String**.
8. In the **Length** field, type the data length of the diameter AVP.
9. In the **Vendor ID** field, type the vendor ID of the diameter Vendor Specific Attribute (VSA).
10. Click **Finished**.

You have configured Diameter AVP.

Updating Signatures for Application Recognition

Overview: Updating classification signatures

Classification signatures define different types of traffic that Policy Enforcement Manager™ (PEM) can recognize, through Traffic Intelligence™. PEM™ recognizes a predefined set of signatures for common applications and application categories that are updated periodically. You can download signature updates from F5 Networks, and schedule the system to automatically update the signatures. You can also manually install the classification signatures and updates, for example, if the BIG-IP® system does not have Internet access.

Task Summary

Scheduling automatic signature updates

Scheduling automatic signature updates

You can set up the BIG-IP® system to automatically update the classification signatures. This ensures that the system always has the latest classification signature files.

1. On the Main tab, click **Traffic Intelligence > Classification > Signature Update**.
The Signatures screen opens.
2. Click **Check for Updates** to manually upload a signature file update if one is available.
You see the current date and time in the **Latest Update Check** setting of the Signature Definitions area.
3. To upload a signature file update, in the Signature Definitions area, click **Import Signatures**.
The Applications screen displays a **Signatures File** field where you can select the new signature file.
4. In the **Signatures File** field, click **Choose File** to navigate to the previously uploaded signatures file.
5. Click **Upload**.
A message displays indicating whether your upload was successful.
6. For the **Automatic Updates Settings**, in the **Signature Update** screen, select **Enabled**.
7. From the **Update Schedule** setting, select **Daily**, **Weekly**, or **Monthly** to specify how often you want the system to check for updates.
8. Click **Update** to save your settings.

The signature updates take effect immediately.

Creating Custom Classifications

Overview: Creating custom classifications

Traffic Intelligence analyzes and identifies higher level protocols and applications. It has the ability to detect applications and protocols in Service Provider networks, for example, HTTP, popular P2P, and top categories (Audio/Video, File Transfer, Instant Messaging, Mail, P2P, Web). It provides an application update mechanism, which in turn, provides the ability to keep up with new, modified, or obsolete applications without going through software release upgrades. IP traffic classifications are based on the IP protocol field of the IP header (IANA protocol).

Note: You can update the library (so) and signature definitions for web traffic (cpm) with hitless upgrade in Policy Enforcement Manager™ (PEM™).

Task summary

Determining and adjusting traffic classifications

Creating a category

Creating classification presets

Creating a custom URL database

Using iRules with classification categories and applications

Modifying iRule event for URL categories

Determining and adjusting traffic classifications

The BIG-IP® system classifies many categories of traffic and specific applications within those categories. You can determine which categories and applications of traffic the system can classify, and find out information about them such as their application or category ID.

1. On the Main tab, click **Traffic Intelligence > Classification > Applications**.
The Applications screen displays a list of the supported classification categories.
2. To view the applications in each category, click the + icon next to the category.
3. To view or edit the properties of the application or category, click the name to open its properties screen.

Tip: Here you can view the application or category ID number.

4. Click **Update** to save any changes.

Creating a category

On the BIG-IP® system, you can create customized categories for classifying traffic if the predefined categories are not sufficient for your needs. For example, if you plan to create new application types unique to your organization, you can create a category to group them together.

1. On the Main tab, click **Traffic Intelligence > Classification > Applications**.
The Applications screen displays a list of the supported classification categories.
2. Click **Create**.
The New Application screen opens.
3. From the **Type** list, select **Category**.
4. In the **Name** field, type a name for the classification category.
5. In the **Description** field, type optional descriptive text for the classification presets.
6. In the **Category ID** field, type an identifier for this category, a unique number.
7. For the **Application List** setting, move applications that you want to associate with this category from the **Unknown** list to the **Selected** list.
If the applications are not listed yet, you can associate the applications with the category when you create them.
8. Click **Finished**.

You have created custom applications to handle traffic.

Creating classification presets

On the BIG-IP® system, you can create classification preset settings for a classification policy that you have previously created.

1. On the Main tab, click **Traffic Intelligence > Classification > Presets**.
The Presets screen displays a list of the supported classification categories.
2. Click **Create**.
The New Presets screen opens.
3. In the **Name** field, type a name for the application.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. For the **Policy** setting, move the classification policies from **Available** list to the **Selected** list, to create a new preset.
6. In the **Allow Reclassification** list, **Enabled** is the default selection.
7. In the **Flow Bundling** list, **Enabled** is the default selection.
8. In the **Cache Results** list, **Enabled** is the default selection.
9. Click **Finished**.

Overview: Creating custom classifications

Creating a category

Creating a custom URL database

Creating a custom URL database

You can create a customized URL database that can be used for adding custom URLs and categories.

1. On the Main tab, click **Traffic Intelligence > URL Categorization > Feed Lists**.
The URL DB feed list screen opens.
2. Click **Create**.
The New Feed List screen opens.
3. In the **Name** field, type a unique name for the URL feed list.

4. In the **Description** field, type optional descriptive text for the URL feed list.
5. In the **Category ID** field, select a category name from the drop-down list.
6. In the URL DB Location area, select the appropriate option for URL DB location.

Option	Description
--------	-------------

File	Click the Browse button, and select the <code>customdb</code> file. The <code>customdb</code> file should be present on your machine, and is not present on the BIG-IP system. The <code>customdb</code> file is a CSV file of the format. The format is: URL/IPv4 [,cat1] [,cat2]...
-------------	--

Note: The non-IP URL should have a IANA registered top level domain. The URL category ID should be in the form of integer and the range is 24576 to 32767.

For example, sample lines of `customdb` entry is:

```
weather.gov, 28678
pconline.com.cn, 28679
kannadaprabha.com, 28680
yandex.ru, 28677, 28676, 28681
pitt.edu, 28682
```

FTP	Type the ftp location and the User and Password .
HTTP	Type the HTTP location and the User and Password .
HTTPS	Type the HTTPS location and the User and Password .

7. In the **Poll Interval** field, type the time interval in seconds at which the url needs to be polled.
8. On the Main tab, click **Traffic Intelligence > URL Categorization > Policies**.
The Policy list opens.
9. Click **Create**.
The New Policy screen opens.
10. In the **Name** field, type a unique name for the URL category policy.
11. In the **Description** field, type optional descriptive text for the URL category policy.
12. In the Feed List area, select the feed list that you created to attach to the policy.
13. Click **Finished**.

The category lookup is done in the custom database, and the URL list is loaded into the custom database through file input. You can also perform URL categorization by looking up the server name indication (SNI) in SSL traffic.

Using iRules with classification categories and applications

If you are using custom classification categories or applications, you can use iRules® to identify the traffic for the custom classifications, or you can initiate an action based on how the traffic is classified.

1. On the Main tab, click **Local Traffic > iRules**.
2. Click **Create**.
3. In the **Name** field, type a 1- to 31-character name.
4. In the **Definition** field, type the syntax for the iRule using Tool Command Language (Tcl) syntax.

For example, to classify traffic as `xxx_app`, a custom classification application that you created, you can use this iRule:

```
when HTTP_REQUEST {
    if { [HTTP::header "Host"] contains "xxx" } {
        CLASSIFY::application set xxx_app
    }
}
```

For example, to perform an action (in this case, drop) on traffic classified as `xxx_app`, you can use this iRule:

```
when CLASSIFICATION_DETECTED {
    if { [CLASSIFICATION::APP == "xxx_app"]} {
        drop
    }
}
```

For complete and detailed information about iRules syntax, see the F5 Networks DevCentral web site <http://devcentral.f5.com>.

5. Click **Finished**.

After creating the iRules, you must assign them as resources for each relevant virtual server on the BIG-IP® system.

Modifying iRule event for URL categories

On the BIG-IP® system, you can modify iRules® Event settings for URL categories.

1. On the Main tab, click **Traffic Intelligence > URL Categorization > Categories**.
2. Select a URL category.
The URL Properties screen opens.
3. In the **Name** field, type a unique name for the URL category policy.
4. In the **Description** field, type optional descriptive text for the classification presets.
5. In the **Category ID** field, type an identifier for this category, a unique number.
6. In the **iRule Event** field, select the appropriate setting.
 - To trigger an iRule event for this category of traffic, select **Enabled**. You can then create an iRule that performs an action on this type of traffic.
 - If you do not need to trigger an iRule event for this category of traffic, select **Disabled**.

Note: *CLASSIFICATION::DETECTED is the only event that is supported.*

You have modified an iRule event setting for an existing URL category.

Classification iRule commands

When the BIG-IP® system identifies a specific type of traffic with iRules® enabled, it triggers a `CLASSIFICATION_DETECTED` event. You can use the commands within iRules for additional system flexibility to classify the flow as one or more of the application or category classifications. The `CLASSIFY` commands are available from the `HTTP_REQUEST` or `HTTP_RESPONSE` iRule events.

iRule Command	Description
<code>CLASSIFICATION::app</code>	Gets the name of the classified application (the most explicit classified application).
<code>CLASSIFICATION::category</code>	Gets the category of the application.
<code>CLASSIFICATION::disable</code>	Disables the classification for a flow.
<code>CLASSIFICATION::enable</code>	Enables the classification for a flow.
<code>CLASSIFICATION::protocol</code>	Gets the name of the classified protocol (the least explicit classified application).
<code>CLASSIFY::application set <i>appname</i></code>	Classifies the flow as <i>appname</i> and associates the category that <i>appname</i> belongs to.
<code>CLASSIFY::category set <i>catname</i></code>	Classifies the flow as <i>catname</i> and also associates the flow with the unknown category.
<code>CLASSIFY::application add <i>appname</i></code>	Adds the application <i>appname</i> to the classification statistics.
<code>CLASSIFY::category add <i>catname</i></code>	Adds the category <i>catname</i> to the classification statistics.

Configuring PEM with Local Traffic Policies

Overview: Creating local traffic policy rules for PEM

When you use Policy Enforcement Manager™ (PEM™), you can create a policy and attach it to traffic policy presets (ce_pem). In the LTM profiles classification (classification_pem), the preset should be ce_pem. The virtual server should have classification profile and SPM profile.

Local traffic policies can include multiple rules. Each rule defines the signature and consists of a condition. Actions are to be performed if the condition holds. Multiple signatures can be assigned to one policy, so you can create a local traffic policy that works with PEM and includes multiple rules that do different things depending on the conditions you set up. In this type of CE policy, each rule can include an application or category or both. The application and category can either be custom or defined applications and categories.

Task Summary

About strategies for local traffic policy matching

About creating custom local traffic policy rules for CE profile

Creating custom local traffic policy for PEM

Creating custom local traffic policy rules for PEM

Creating a virtual server for SSL traffic policy enforcement

Associating a published local traffic policy with a virtual server

Creating custom local traffic policy for PEM

About strategies for local traffic policy matching

Each BIG-IP® local traffic policy requires a matching strategy to determine which rule applies if more than one rule matches.

The BIG-IP local traffic policies provide three predefined policy matching strategies: a first-match, best-match, and all-match strategy. Each policy matching strategy prioritizes rules according to the rule's position within the Rules list.

As needed, you can create a user-defined best-match strategy to customize the precedence (order of preference) of added operands and selectors. For example, to meet your preferred operand and selector combinations, you might create a user-defined best-match strategy that changes the precedence of added operands and selectors, compared to the predefined best-match strategy.

Note: In a best-match or first-match strategy, a rule without conditions becomes the default rule, when the rule is the last entry in the Rules list.

Table 1: Policy matching strategies

Matching strategy	Description
all-match strategy	<p>An <i>all-match strategy</i> starts the actions for all rules in the Rules list that match.</p> <hr/> <p>Note: In an <i>all-match strategy</i>, when multiple rules match, but specify conflicting actions, only the action of the best-match rule is implemented. A best-match rule can be the lowest ordinal, the highest priority, or the first rule that matches in the Rules list.</p> <hr/>
best-match strategy	<p>A <i>best-match strategy</i> selects and starts the actions of the rule in the Rules list with the best match, as determined by the following factors.</p> <ol style="list-style-type: none"> 1. A best-match strategy selects the rule with the most conditions, ignoring details about the conditions. 2. If a rule with the most conditions is not determined, then the best-match strategy selects the rule with the highest priority condition types. The best-match strategy sorts the condition types, highest priority first, comparing one at a time until a higher priority is found. For example, a priority sequence of 0,1,3,4,6 wins over 0,1,3,5,7 because 4 is a higher priority than 5. 3. If a rule with the highest priority condition types is not determined, then the best-match strategy selects the rule with equal match types over other match types, such as starts-with, ends-with, or contains, and processes according to condition type priority. 4. If a rule of equal match types is not determined, then the best-match strategy uses an ordinal (the precedence of the operand). <hr/> <p>Note: In a <i>best-match strategy</i>, when multiple rules match and specify an action, conflicting or otherwise, only the action of the best-match rule is implemented. A best-match rule can be the lowest ordinal, the highest priority, or the first rule that matches in the Rules list.</p> <hr/>
first-match strategy	<p>A <i>first-match strategy</i> starts the actions for the first rule in the Rules list that matches.</p>

[Overview: Creating local traffic policy rules for PEM](#)

[Overview: Creating local traffic policy rules for PEM](#)

[About creating custom local traffic policy rules for CE profile](#)

About creating custom local traffic policy rules for CE profile

Classification signatures are added as rules in the local traffic policy. The classification signatures can be used for many standard categories and applications. In addition, you can create custom categories and applications. When you use Policy Enforcement Manager™ (PEM™), you can create a policy and attach it to traffic policy presets (ce_pem). In the LTM profiles classification (classification_pem), the preset should be ce_pem. The virtual server should have classification profile and SPM profile.

Local traffic policies can include multiple rules. Each rule defines the signature and consists of a condition. Actions are to be performed if the condition holds. Multiple signatures can be assigned to one policy, so you can create a local traffic policy that works with PEM and includes multiple rules that do different things depending on the conditions you set up. In this type of CE policy, each rule can include an application or category or both. The application and category can either be custom or defined applications and categories.

Task Summary

Overview: Creating local traffic policy rules for PEM

About strategies for local traffic policy matching

Creating custom local traffic policy for PEM

Creating custom local traffic policy for PEM

Before you modify rules on existing policies, you must set up an application or category (**Traffic Intelligence > Classification**).

You can add rules to define conditions and run specific actions for different types of application traffic in Policy Enforcement Manager™ (PEM™). For example, if you create an application signature for company A and want to send traffic from company A's website, you can perform actions, such as bandwidth control and disable **Gate status** from PEM. This is a rule that can be assigned to an existing policy.

1. On the Main tab, click **Local Traffic > Policies**.

For more information about local traffic policies, refer to *BIG-IP® Local Traffic Manager™ Implementations*.

The Policy List screen opens.

2. Click **create**.

The New Policy List screen opens.

3. In the **Policy Name** field, type a unique name for the policy, for example `companyA`.

4. In the **Description** field, type descriptive text that identifies the policy definition.

5. From the **Strategy** list, select the action that is executed when there are multiple rules that match.

Rule	Description
All	Uses the first or best strategy to resolve the conflict of rule match.
Best	Applies the actions of the rule specified in the list of defined strategies for the associated policy.
First	Applies the actions of only the first rule. This implies that the rule with the lowest ordinal, highest priority or first in the list is executed.

6. From the **Type** list, select the **Traffic Policy** to create a custom signature.

7. Click **Create Policy** to create a policy that manages traffic assigned to a virtual server.

8. Click the down arrow for Save Draft. Select **Save Draft Policy** to save the policy as a draft or **Save and Publish policy** to publish a policy and assign it to a virtual server.

You should be able to create a rule for the Draft Policies list.

9. Click the name of the draft policy you just created.

The Draft Policy screen opens.

10. From the **Rules** list, select **Create**.

The New Rule screen opens.

11. In the **Name** field, type a unique name for the rule.

12. In the **Description** field, type descriptive text that identifies the rule definition.

13. In **Match all of the following conditions**, click + and specify the conditions.

For example, select **Client SSL**, **cipher**, **contains** and type `COMPAT:AES128-GCM-SHA256`, **request**

14. Click **Add**.

15. In **Do the following when the traffic is matched**, click + and specify the actions:

For example, select **Enable**, **cache**, at **request**.

16. Click **Save**.

Now you have added a new rule to the existing policy. When you send traffic that matches the rule you defined, you should be able to see the application or category you have configured.

Creating custom local traffic policy rules for PEM

You can create a new strategy for your policy in Policy Enforcement Manager™ (PEM™).

1. On the Main tab, click **Local Traffic > Policies > Strategy List**.
The Strategy List screen opens.
2. Click **Create**.
The New Strategy List screen opens.
3. In the **Name** field, type a unique name for the strategy definition.
4. In the Operands area, define the application traffic to which this rule applies. Specify these values and use default values for the remainder.
 - a) From the **Operand** list, select **http-host**.
 - b) From the **Event** list, select **request**.
 - c) From the **Selector** list, select **all**.
 - d) From the **Condition** list, select **ends-with**.
 - e) Type the value; for example, `f5.com`.
 - f) Click **Finished**.

Now you have created a strategy list and changed how the system processes the operands by reordering the list of definitions.

Creating a virtual server for SSL traffic policy enforcement

The BIG-IP® system allows SSL pass through mode to collect certificate information. You have to define a virtual server that references SSL pool and classifies SSL traffic for policy enforcement.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the **Create** button.
The New Virtual Server screen opens.
3. In the **Name** field, type a unique name for the virtual server.
4. For a network, in the **Destination Address** field, type an IPv4 or IPv6 address in CIDR format to allow all traffic to be translated.

The supported format is address/prefix, where the prefix length is in bits. For example, an IPv4 address/prefix is `0.0.0.0/0`, and an IPv6 address/prefix is `::/0`.
5. In the **Service Port** field, type `443` or select **HTTPS** from the list.
6. From the **Configuration** list, select **Advanced**.
7. From the Classification list, select **Enabled**, for the BIG-IP system to enable classification for virtual servers when a policy enforcement listener is created.
8. From the **Policy Enforcement Profile** list, select the name of the policy enforcement profile that you previously created.
9. Click **Finished**.

10. From the **Default Persistence Profile** list, select **ssl**.
This implements simple persistence, using the default ssl profile.
11. In the Policies area, click the **Manage** button.
12. For the **Policies** setting, from the **Available** list, select the name of the iRule that you want to assign, and use the buttons to move the name into the **Enabled** list.

You have created a virtual server for SSL traffic. The virtual server that references SSL pools appears in the Virtual Servers list.

Associating a published local traffic policy with a virtual server

After you publish a local traffic policy, you associate that published policy with the virtual server created to handle application traffic.

1. On the Main tab, click **Local Traffic > Virtual Servers**.
The Virtual Server List screen opens.
2. Click the name of the virtual server you want to modify.
3. On the menu bar, click **Resources**.
4. In the Policies area, click the **Manage** button.
5. For the **Policies** setting, from the **Available** list, select the local traffic policy you previously created, and move it to the **Enabled** list.
6. Click **Finished**.

The published policy is associated with the virtual server.

Overview: Creating local traffic policy rules for PEM

Creating a virtual server for SSL traffic policy enforcement

Creating custom local traffic policy for PEM

Creating custom local traffic policy for PEM

Before you modify rules on existing policies, you must set up an application or category (**Traffic Intelligence > Classification**).

You can add rules to define conditions and run specific actions for different types of application traffic in Policy Enforcement Manager™ (PEM™). For example, if you create an application signature for company A and want to send traffic from company A's website, you can perform actions, such as bandwidth control and disable **Gate status** from PEM. This is a rule that can be assigned to an existing policy.

1. On the Main tab, click **Local Traffic > Policies**.
For more information about local traffic policies, refer to *BIG-IP® Local Traffic Manager™: Implementations*.
The Policy List screen opens.
2. Click **create**.
The New Policy List screen opens.
3. In the **Policy Name** field, type a unique name for the policy, for example `companyA`.
4. In the **Description** field, type descriptive text that identifies the policy definition.
5. From the **Strategy** list, select the action that is executed when there are multiple rules that match.

Rule	Description
All	Uses the first or best strategy to resolve the conflict of rule match.
Best	Applies the actions of the rule specified in the list of defined strategies for the associated policy.
First	Applies the actions of only the first rule. This implies that the rule with the lowest ordinal, highest priority or first in the list is executed.

6. From the **Type** list, select the **CE Profile** to create a custom signature.
7. Click **Create Policy** to create a policy that manages traffic assigned to a virtual server.
8. Click the down arrow for Save Draft. Select **Save Draft Policy** to save the policy as a draft or **Save and Publish policy** to publish a policy and assign it to a virtual server.
You should be able to create a rule for the Draft Policies list.
9. Click the name of the draft policy you just created.
The Draft Policy screen opens.
10. From the **Rules** list, select **Create**.
The New Rule screen opens.
11. In the **Name** field, type a unique name for the rule.
12. In the **Description** field, type descriptive text that identifies the rule definition.
13. In **Match all of the following conditions**, click + and specify the conditions.
For example, select **Client SSL**, **cipher**, **contains** and type `COMPAT:AES128-GCM-SHA256`, **request**
14. Click **Add**.
15. In **Do the following when the traffic is matched**, click + and specify the actions:
For example, select **Enable**, **cache**, at **request**.
16. Click **Save**.

Now you have added a new rule to the existing policy. When you send traffic that matches the rule you defined, you should be able to see the application or category you have configured.

Overview: Creating local traffic policy rules for PEM

Associating a published local traffic policy with a virtual server

Overview: Configuring policy and RADIUS updates

Configuring Policy and RADIUS Updates

Overview: Configuring policy and RADIUS updates

Policy Enforcement Manager™ (PEM™) enables you to schedule policy reevaluations and radius updates on the BIG-IP® system in the following two ways:

- You can configure the interval for reevaluation of policies, for a subscriber session, by configuring the re-evaluation interval. The BIG-IP system evaluates changes in the policy for traffic, once the re-evaluation interval is configured.
- The RADIUS traffic contains the subscriber and IP address information that is monitored by the BIG-IP system. If you enable the timeout interval, the BIG-IP system avoids repeated deletion and creation of the subscriber during the configured interval rate.

Configuring PEM options

You can set up the BIG-IP® system to schedule an interval that sets policy reevaluation and RADIUS re-transmission updates periodically.

1. On the Main tab, click **Policy Enforcement > Options**.
The Options screen opens.
2. In the Policy Options area, specify (in seconds) the **Policy Re-evaluation Interval** at which the policy re-evaluation is triggered, to evaluate the flow policy again.
The re-valuation interval is only for active flows.
For example, a subscriber is provisioned over Gx which has a policy to allow Netflix with some bandwidth. The subscriber is able to watch a movie using the Netflix service. However, consider that the PCRF installs a policy for this subscriber to block Netflix over the Gx interface. Then, while the subscriber is viewing the content, the Netflix content is blocked for the subscriber after the configured re-evaluation interval.
3. In the RADIUS Options area, for the **Re-Transmit Timeout** setting, select **Enabled** and specify the time in seconds. If you select **Disabled**, each RADIUS message is handled as a new message and this might lead to deletion and creation of sessions even though the radius message is a duplicate. This is the timeout after which the RADIUS message is considered as a new message, by the BIG-IP system.
4. In the RADIUS Options area, for the **Clear Sessions upon NAS Reboot** setting, select **Enabled** to remove all the PEM sessions that are associated with the NAS-IP-Address received in the RADIUS Acct-ON or Acc-OFF request packet.
5. In the Quota Management Options area, for the **Default Rating Group** setting, select **Create** to create a new rating group for quota management.
This takes you to the **Policy Enforcement > Rating Groups > New Rating Group** screen. Click **Policy Enforcement > Options** to go back to options screen.
6. In the Statistics Options area, for the **Analytics Mode** setting, select **Enabled** to use analytics reporting. Select the external logging such as HSL endpoint in the **External Log Publisher** setting.
This generates Application Visibility and Reporting (AVR) PEM reports, in a timely manner through graphs.
7. From the **Subscriber Aware** list, select **Enabled** to display the statistics per subscriber.

This generates Application Visibility and Reporting (AVR) PEM reports, in a timely manner through graphs.

8. For the **Content Insertion Options** setting, in the **Throttling** fields, type the time used to set the maximum wait time before Policy Enforcement Manager™ applies the insert action again on the same subscriber.

***Note:** The insert actions do not conflict with each other.*

The policy and RADIUS updates take effect immediately.

Creating Custom TAC DB

Overview: Custom TAC DB

In Policy Enforcement Manager™ (PEM™), custom Type Allocation Code DB (TAC DB) is either the one delivered with the release, or it can be downloaded from a configurable location. The activation is hitless and does not require system restart. The use of the TAC DB is licensed by default with PEM license, as well as the one the user can configure. This manages the licensed TAC DB. You are now able to easily update TAC DB automatically whenever a new version is available. In PEM, you can also provide a location for your own TAC DB. The BIG-IP® system will download the custom DB, and initial updates from the configured location.

Task summary

Creating a custom TAC DB

Creating a custom TAC DB

You can use the customdb component in PEM™ to define a custom TAC DB.

1. On the Main tab, click **Policy Enforcement > Global Options > TAC DB**.
The TAC DB screen opens.
2. In the Custom TAC DB Configuration area, for the **Priority** setting, select **High** or **Low** from the drop-down menu to specify that if the Custom TAC DB is built-in.
3. In the TAC DB setting, for the **Name** field, type a name for the custom TAC DB.
4. In the TAC DB setting, for the **Upload New From** drop-down menu, select **File** to choose a file with TAC code information that is located on your client machine or **Location** to set the URL location of the TAC DB. The URL can be, HTTP, HTTPS, FTP, SFTP or file on BIG-IP file system.
If you choose **Location**, additional fields in TAC DB setting appear.
5. In the TAC DB setting, for the **Description** field, type the description for this custom TAC DB.
6. In the TAC DB setting, for the **TAC DB Location**, type the URL location where the TAC DB file is stored.
7. In the TAC DB setting, for the **Username** field, type the username for access to the custom TAC DB information.
8. In the TAC DB setting, for the **Password** field, type the password to access the custom TAC DB information.
9. In the Sync Mode setting, select **Automatic** to automatically upload files from the location and specify how often the polling interval can be configured in weeks. Select **Manual** to manually synchronize the files.
10. Click **Update** to save any changes.

Overview: Custom TAC DB

Enforcing Policy and Classification on IP Protocols

About enforcing policy and classification on IP protocols

The BIG-IP® system now provides classification and policy enforcement on all non-TCP and non-UDP traffic, which includes IPsec traffic. The Policy Enforcement Manager™ is able to classify and enforce any action on virtually any type of IP traffic. This enables detection of IPsec, ICMP, GRE, and other IP protocols (especially tunneling) for the service providers. For IPsec, Encapsulating Security Payloads (ESP) and Authentication Headers (AH) protocols are used, in both tunnel and transport modes.

A bottom hudfilter forwards non-TCP and non-UDP traffic for both classification and policy enforcement.

Note: HTTP redirect is not supported. Based on the protocol, not all actions work and some traffic is not steered.

Important: You can use SNAT, only when you forward ICMP and ICMPv6 traffic.

Creating Any IP profiles for PEM

Before you create multiple Any IP profiles, you must create a listener in Policy Enforcement Manager™ (PEM™), which creates a virtual server with Any IP profile.

You can create a new Any IP profile through local traffic management in PEM.

1. On the Main tab, click **Local Traffic > Profiles > Protocol > Any IP**.
The Any IP screen opens.
2. Click **Create**.
The New Any IP Profile screen opens.
3. From the **Parent Profile** list, select the default **ipother** or any other Any IP profile, from where the new profile can inherit the settings.

Note:

You will see multiple Any IP profiles in the list only if you have created the profiles earlier.

4. To specify the idle timeout, click **Custom**, select **Specify**, and type a value (in seconds). The idle time out specifies the number of seconds for which a connection is idle before the connection is eligible for deletion.
5. Click **Finished**.

Now you have created a new Any IP profile. You can view non-TCP and UDP traffic that passes through the BIG-IP system (**Statistics > Classification > Statistics**).

Updating Any IP profile

If you have created other Any IP profile and you want to attach this profile to the Any IP traffic, then you can attach the profile through local traffic management in Policy Enforcement Manager™ (PEM™).

1. On the Main tab, click **Local Traffic > Virtual Servers > Virtual Server List**.
The Virtual Server List screen opens.
2. Select any virtual server.
The virtual server properties screen opens.
3. From the **Protocol** list, select ***All Protocols**.
Any IP profile settings displays.
4. From the **Any IP Profile** list, select the default setting **ipother**, or any other Any IP profile from where the new profile can inherit the settings.

Note:

You will see multiple Any IP profiles from the list only if you have created the profiles earlier.

5. Click **Update**.

Now you have updated the Any IP profile and attached it to the Any IP traffic.

IPOther filter for current PEM actions

The policy actions configured in the Policy Enforcement Manager™ can support non-TCP and non-UDP traffic flows. This table contains the information that highlights the actions supported for non-TCP and non-UDP traffic.

Action	All non-TCP and non-UDP flows
Forwarding	Only non-tunnel protocols. <i>Note: ICMP traffic can be steered.</i>
Service-chain	Only non-tunnel protocols. <i>Note: ICMP traffic can be steered.</i>
Cloning	Yes
BWC (both directions)	Yes
L2 QoS markings (both directions)	Yes
Flow Reporting	Yes
Session Reporting	Yes
Gate status drop	Yes
Quota	Yes
HTTP-redirect	No
Modify HTTP headers	No
iRules	CLIENT_DATA and CLIENT_ACCEPTED iRules only (like UDP filter).

Detecting Tethering, Device Operation System and Type

Overview: Detecting Tethering, Device Operating System, and Type

You can now gather insight on the device usage of subscribers, with the detection of Device Type and Operating System (DTOS). The patterns can be used for business intelligence gathering, as well as for customizing subscriber plans based on their usage. The Policy Enforcement Manager™ (PEM™) provides the ability to report various types of subscriber information and application visibility to external analytic systems through syslog and other IPFIX reporting methods. The device type is identified based on the IMEI number, of which the first eight numbers are from the Type Allocation Code (TAC) number. The operating system (OS) detection is done by user-agent parsing, TCP/IP fingerprinting, and by looking up TAC code in the TAC database. Furthermore, reporting can be configured for DTOS and tethering action. A report (optional) is sent the first time PEM retrieves device OS information, and if a change is detected in the OS name, a report is sent again to the configured destination.

When tethering is enabled, details of the state change are sent in a report. When DTOS is enabled, the details of the state changes when there is a change in TCP fingerprinting or OS value defined by the user agent. Also, a report (HSL log) is sent out whenever there is a change in TCP fingerprinting, TCP OS, or user agent OS values. The default sampling interval for DTOS is at 10 flows. The tethering sampling interval is 180 seconds.

Task summary

Configuring device type, OS, and tethering

Configuring PEM policy action with tethering

Creating a high-speed logging rule for device detection and tethering

Configuring device type, OS, and tethering

You can configure the PEM™ policy to look up the device name and details (for monitoring purposes) and enable tethering detection.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

Tip: All rules in a policy are run concurrently. Precedence takes effect when there are conflicting rules. The conflict occurs when the traffic matches two rules and the policy actions from these rules differ. For example, if you have rule 1 with precedence 10 and **Gate Status** disabled for a search engine, and you have rule 2 with precedence 11 and **Gate Status** enabled, then rule 1 is processed first because it

has higher precedence. Rules conflict if they have identical or overlapping classification criteria (for the traffic that matches more than one rule). In some cases, different policy actions are not conflicting, and hence, applied in parallel.

6. From the **Device and Tethering Detection** list, in the **Device Type OS Detection** setting, select **Enabled**.

***Note:** If you enable device detection, all the filters are disabled for the policy rule.*

***Note:** When the custom TACDB file is generated, it is stored at the location `/var/local/pem/dtos/`.*

7. From the **Device and Tethering Detection** list, in the **Tethering Detection** setting, select **Enabled**.

***Note:** If you enable tethering, classification is disabled for the policy rule.*

8. Click **Finished**.

You have created a rule that applies to policy based on device OS, type, and tethering.

Configuring PEM policy action with tethering

You can also attach a policy action after detecting tethering. The tethering detected state can be used as a TCL filter in a custom filter for a Policy Enforcement Manager™ (PEM™) rule.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the high precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.

***Note:** TCL filter creation action should have high precedence.*

6. Click the Custom tab.
The Custom Criteria setting opens.
7. In the **iRule Expression** field, specify the TCL syntax that defines a custom iRule action, which can be later attached to a policy enforcement rule.

```
PEM::session info tethering detected [IP::client_addr] == {1}
```

The expression in the example evaluates to true, if a subscriber is tethering. You can select the **Wrap Text** check box to wrap the definition text, and select the **Extend Text Area** check box to increase the field space of format scripts.

8. From the **Gate Status** list, select **Disable**, to block the traffic for a subscriber who is tethering.

***Note:** If you disable **Gate Status**, the traffic is blocked.*

9. To apply bandwidth policy, for rate control to downlink traffic, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

***Note:** You can assign any previously created static or dynamic bandwidth control policies. However, F5® does not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create to enforce dynamic QoS settings provisioned by the PCRF.*

Depending on the bandwidth control policy, PEM™ restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, or any combination of these.

10. To apply bandwidth policy, for rate control to uplink traffic and per category of application, in the **Bandwidth Controller** setting, select the name of a bandwidth control policy.

***Note:** You can assign any previously created static or dynamic bandwidth control policies. However, we do not recommend using the **default-bwc-policy**, which the system provides, nor the **dynamic_spm_bwc_policy**, which you can create for communicating with the PCRF.*

Depending on the bandwidth control policy, PEM restricts bandwidth usage per subscriber, group of subscribers, per application, per network egress link, per category of applications or any combination of these.

11. Click **Finished**.

You have created a rule that applies to tethering.

Creating a high-speed logging rule for device detection and tethering

You can specify a reporting destination where reports are sent out whenever the subscribers go from a non-tethering state to a tethering state, or vice-versa. Before you can create a high-speed logging (HSL) rule, you need to create a publisher that defines the destination server or pool where the HSL logs are sent. In an enforcement policy, a rule can specify that tethering details are sent to an external high-speed logging server.

1. On the Main tab, click **Policy Enforcement > Policies**.
The Policies screen opens.
2. Click the name of the enforcement policy you want to add rules to.
The properties screen for the policy opens.
3. In the Policy Rules area, click **Add**.
The New Rule screen opens.
4. In the **Name** field, type a name for the rule.
5. In the **Precedence** field, type an integer that indicates the high precedence for the rule in relation to the other rules. Number 1 has the highest precedence. Rules with higher precedence are evaluated before other rules with lower precedence.
6. In the **Reporting** setting, specify where to send the tethering detection data:
 - From the **HSL** list, select the name of the publisher that specifies the server or pool of remote HSL servers to send the logs.
 - From the **Format Script** list, select the format script of the report from the **Format Script** list.

***Note:** The format script is previously configured in **Policy Enforcement > Reporting > Format Script** page.*

7. Click **Finished**.

You have created a rule that sends device detection and tethering data about the traffic to external high-speed logging servers.

Implementation result

The BIG-IP[®] system allows improved insight to subscriber usage with detection of tethering, and device OS and type.

Troubleshooting

PEM troubleshooting

Follow these general troubleshooting suggestions when using Policy Enforcement Manager™ (PEM™):

- If enforcement policies are not enforced as expected, on the VLAN screen for all VLANs set up to receive incoming subscriber traffic, verify that you set **CMP Hash** to **Source Address**.
- If static subscriber policies are not enforced as expected, verify whether you enforced any global, high precedence policies with conflicting actions.
- When sending traffic without RADIUS, the unknown subscriber policy (if specified) is assigned to the first flows from dynamic or static subscribers. Subscriber policies are applied to subsequent flows.

***Note:** An unknown subscriber policy needs to be specified, if there is at least one dynamically provisioned subscriber.*

- Policy changes are applied to new and existing flows within a reasonable time.
- For applications with connections initiated from the Internet (FTP, RTSP, TFTP), the BIG-IP® system needs to have **CMP Hash** set to **Destination Address** on the Internet VLAN. In this case, the end-to-end IP addresses have to be preserved; therefore, SNAT should be disabled on all the virtual servers that the applications will use.
- When importing static subscribers, the file is uploaded in chunks of 1000 subscribers. The system performs a validation check on each chunk. If a validation fails, the subscribers in the current chunk and subsequent chunks are not imported. However, the subscribers loaded in previous chunks are imported onto the system.

***Note:** PEM™ can use 3rd party database, custom DB or iRule for URL categorisation. The onbox 3rd party database is limited to the 20M most used URL and is updated regularly.*

Steering troubleshooting

- In case of service chains (w-steering), set **CMP Hash** to **Source Address** on all the VLANs for which the w-steering action is to be applied.
- For response-side classification, steering, w-steering, and cloning actions are applied after the results (based on destination IP address and port) are cached in the classification database (srdb). Actions are not applied for the first six flows, by default. (This behavior is configurable by the DB variable `tmm.pem.srdb.entry.step`.)

RADIUS troubleshooting

- If static subscribers are not working as expected with RADIUS, check whether you selected the same **Subscriber ID Type** in the `radiusLB` profile (**Local Traffic > Profiles > Services > RADIUS**) as that assigned when creating the static subscriber. (**IMSI** in the static subscriber corresponds to **3GPP IMSI** in the RADIUS profile; **E164** to **Calling Station ID**, and **NAI** to **User Name**.)
- The RADIUS message also needs to specify the same **Subscriber ID Type** as the RADIUS profile. So make sure that if you select **IMSI**, the IMSI number exists in the RADIUS message. This also applies to the **user-name** for NAI, and **calling station-id** for E164.

Gx interface to PCRF troubleshooting

- If you change the IP address of the Gx server in the listener, the change takes effect after you restart TMM using the command: `bigstart restart tmm`.
- For Gx usage monitoring, the threshold is defined on the Policy and Charging Rules Function (PCRF).

Bandwidth control with PEM troubleshooting

- Do not use dynamic bandwidth control policies in preconfigured enforcement policies (either global or subscriber) when the bit rate is managed by the PCRF through PCC dynamic rules.
- Do not use dynamic bandwidth control policies in global enforcement policies if they are also used in subscriber policies.
- For bandwidth controller to work with PCRF, you need to create a default dynamic bandwidth controller with the name `dynamic_spm_bwc_policy`, with eight categories named `cat1` to `cat8` (all set to 100 percent). You must choose a proper max-rate value for this bandwidth controller (typically, close to network capacity dedicated to subscriber traffic).

Important: This bandwidth controller is intended for internal usage only and should not be used for other purposes.

Active sessions troubleshooting (retrieving subscriber data or BIG-IP system information)

- When the BIG-IP system receives policy information from the PCRF for a subscriber, you can verify the active policies on the subscriber session, the subscriber type (static or dynamic) and view subscriber statistics by checking **Active Sessions (Policy Enforcement > Subscribers > Active Sessions)**.
- If you have a static subscriber without an IP address, no active session is created. The incoming RADIUS message has the IP information for the static subscriber and a session is created based on this. When the radius message arrives, verify both the new session and policy attached to the session.
- You can view subscriber information with multiple IP addresses. Static subscribers can have more than two IP addresses of either IPv4 or IPv6 and up to a maximum of 16. Dynamic subscribers can have one IPv4 IP address and one IPv6 IP address.
- If your subscriber type is dynamically provisioned, then your assigned policy can be based on a predefined PCC rule or dynamic PCC rule.
- For information about uplink and downlink traffic (byte count and flows), check **(Policy Enforcement > Subscribers > Statistics)**.
- You can auto-refresh the subscriber session information for 10 to 300 seconds.
- There is a hold time for new subscriber sessions. To change the provisioning hold time, you can use the sys db variable key: `tmm.pem.session.ip-addr.max`.

iRules® troubleshooting

- While running the script, if the BIG-IP system receives an error, ignore the error and implement the next custom action script. Although this is the default behaviour, it is possible to change it with the sys db variable key: `pem.tcl.action.error.abort`.
- If policy priority, event priority, and the rule precedence is the same, then there is no guarantee of order of execution.
- You can use iRule commands to set accounting report interval, but set the accounting interval larger than the BIG-IP interval configuration for the accounting report interval to be effective.

IPsec troubleshooting

- For IPsec to work with Policy Enforcement Manager™ (PEM™), disable the DB variable `ipsec.lookupspi`.

Subscriber and policies active sessions

You can view session records based on subscriber ID or session IP. Policy Enforcement Manager™ contains the information presented in this table. You can access this is in **Active Sessions (Policy Enforcement > Subscribers > Active Sessions)**.

Field	Description
ID	A unique identifier (up to 64 characters) for the subscriber initiating the session, such as a phone number. The subscriber ID type determines the format.
ID type	The format of the subscriber ID attribute. It can be E.164, IMSI, NAI, or Private (RFC 4006).
Subscriber Type	Specifies a dynamically or statically subscriber.
Calling Station	Radius Attribute Value Pair (AVP) type 31 (3GPP TS 29.061 V9.6.0).
Called Station	Radius Attribute Value Pair (AVP) type 30 (3GPP TS 29.061 V9.6.0).
Tower	Specifies the cell tower where subscriber information goes through.
User Name	Displays the format name name@domain.
IMSI	International Mobile Subscriber Identity. A globally unique code number that identifies a GSM, UMTS, or LTE mobile phone user.
IMSEISV	International Mobile Station Equipment Identity Software Version. A globally unique code number that identifies a GSM, UMTS, LTE, or iDEN mobile phone.
Predefined	Specifies the predetermined policy(ies) assigned to the subscriber.
Dynamic	Specifies the dynamic PCC rule applied.
Statistics	Specifies active session statistical information that includes subscriber and session IP identity attributes, assigned policy, and traffic flow information.

Active sessions statistics

You can view subscriber uplink and downlink traffic information. Policy Enforcement Manager™ contains the information presented in this table.

Field	Description
Data Format	Specifies how the system presents the statistics information. The default is Normalized .
Auto Refresh	Automatically updates the screen information at the interval you specify. For example, if you select 60 seconds from the list, the system updates the displayed screen information every 60 seconds. The default is Disabled . When you specify an automatic-refresh interval, the system presents a Stop button for halting the operation, and counts down the seconds to the next update. Select Disabled to turn off automatic refreshing.
Session IP	Specifies the session IP address. The IP address is in either IPv4 or IPv6 format.
Subscriber ID	Specifies a unique identifier subscriber ID.

Field	Description
Uplink	Specifies traffic volume from the subscriber to network.
Downlink	Specifies traffic volume from the network to subscriber.
Current	Specifies current number of flows.
Maximum	Specifies maximum number of open flows.
Total	Specifies accumulated number of flows ever opened by the subscriber.

Configuring subscriber activity log

You can configure the activity logs of the selected subscribers by subscriber or session activity.

1. On the Main tab, click **Policy Enforcement > Subscribers > Activity Log > Configuration**. The Configuration screen opens.
2. From the **Log Publisher** list, select the log publisher that was created. You can create a log publisher in the system at **System > Logs > Configuration > Log Publishers**.
3. From the **Subscriber Type** list, select **Dynamic** (for dynamic provisioning) or **Static** (for static provisioning) subscriber.
4. In the **Subscriber ID** field, type a unique identifier (up to 64 characters) for the subscriber, such as IMSI.
5. Using the Log Subscriber Activity setting, add each subscriber ID to the log settings.
 - a) Type the **Subscriber ID**.
 - b) Click **Add**.
6. To configure settings of the activity logs by sessions, use the Log Session Activity setting to add each session IP to the log settings.
 - a) Type the **Session IP** address.
 - b) Click **Add**.
7. Click **Update**.
Policy Enforcement Manager™ starts generating the subscriber activity logs for the configured subscribers.

You have configured the activity logs settings. Policy Enforcement Manager applies the log settings you assigned and lists subscriber activity and session information.

IPFIX Templates for PEM Events

Overview: IPFIX templates for PEM events

The IP Flow Information Export (IPFIX) Protocol is a logging mechanism for IP events. This appendix defines the IPFIX Information Elements (IEs) and templates used to log F5[®] Policy Enforcement Manager[™] (PEM[™]) events. An *IE* is the smallest form of useful information in an IPFIX log message, such as an IP address or a timestamp for the event. An *IPFIX template* is an ordered collection of specific IEs used to record one IP event, such as the acceptance of a network packet. In PEM, the IPFIX publisher delivers PEM records at the session, flow, and transaction level.

About IPFIX Information Elements for PEM events

Information Elements (IEs) are individual fields in an IPFIX template. An IPFIX template describes a single Policy Enforcement Manager[™] (PEM[™]) event.

IANA-defined IPFIX Information Elements

IANA maintains a list of standard IPFIX Information Elements (IEs), each with a unique Element Identifier. The F5[®] PEM[™] IPFIX implementation uses a subset of these IEs to publish PEM events. This subset is summarized in the table.

Information Element (IE)	ID	Size (Bytes)
destinationIPv4Address	12	4
destinationIPv6Address	28	16
destinationTransportPort	11	2
ingressVRFID	234	4
protocolIdentifier	4	1
sourceIPv4Address	8	4
sourceIPv6Address	27	16
sourceTransportPort	7	2

IPFIX enterprise Information Elements

IPFIX provides for enterprises to define their own Information Elements. F5[®] currently uses the following non-standard IEs for PEM[™] events:

Information Element (IE)	ID	Size (Bytes)
3gppParameters	12276 - 57	Variable
applicationCategoryId	12276 - 48	2
concurrentFlows	12276 - 59	2
downlinkVolume	12276 - 88	8
durationSec	12276 - 60	2
lastRecordSent	12276 - 63	8
newFlows	12276 - 64	2
observationTimeSeconds	12276 - 90	8
recordReason	12276 - 66	1
recordType	12276 - 54	1
reportId	12276 - 55	4
reportVersion	12276 - 56	Variable
subscriberId	12276 - 71	Variable
subscriberIdType	12276 - 72	Variable
successfulTransactions	12276 - 68	4
terminatedFlows	12276 - 69	2
timestampMsec	12276 - 91	2
totalTransactions	12276 - 73	2
uplinkVolume	12276 - 89	8

Information Element (IE)	ID	Size (Bytes)
applicationCategoryId	12276 - 48	2
downlinkVolume	12276 - 88	8
flowStartMilliseconds	12276 - 50	2
flowStartSeconds	12276 - 51	8
flowStopMilliseconds	12276 - 52	2
flowStopSeconds	12276 - 53	8
observationTimeSeconds	12276 - 90	8
recordType	12276 - 54	1
reportId	12276 - 55	4
reportVersion	12276 - 56	Variable
subscriberId	12276 - 71	Variable
subscriberIdType	12276 - 72	Variable
timestampMsec	12276 - 91	2
totalTransactions	12276 - 73	2
uplinkVolume	12276 - 89	8
urlCategoryId	12276 - 87	2

Information Element (IE)	ID	Size (Bytes)
vlanId	12276 - 92	2

Information Element (IE)	ID	Size (Bytes)
applicationCategoryId	12276 - 48	2
classification	12276 - 49	Variable
downlinkVolume	12276 - 88	8
httpHostname	12276 - 74	Variable
httpHostnameTruncated	12276 - 75	1
httpResponseCode	12276 - 76	2
httpUrl	12276 - 77	Variable
httpUrlTruncated	12276 - 78	1
httpUserAgent	12276 - 79	Variable
httpUserAgentTruncated	12276 - 80	1
recordType	12276 - 54	1
reportId	12276 - 55	4
reportVersion	12276 - 56	Variable
skippedTransactions	12276 - 82	2
subscriberId	12276 - 71	Variable
subscriberIdType	12276 - 72	Variable
transactionNumber	12276 - 81	2
transactionStartMilliseconds	12276 - 83	2
transactionStartSeconds	12276 - 84	8
transactionStopMilliseconds	12276 - 85	2
transactionStopSeconds	12276 - 86	8
uplinkVolume	12276 - 89	8
urlCategoryId	12276 - 87	2
vlanId	12276 - 92	2

Note: IPFIX, unlike NetFlow v9, supports variable-length IEs, where the length is encoded within the field in the Data Record. NetFlow v9 collectors (and their variants) cannot correctly process variable-length IEs, so they are omitted from logs sent to those collector types.

About individual IPFIX templates for each event

F5® uses IPFIX templates to publish PEM™ events.

Session logs

This IPFIX template is used for session records used for HSL reporting.

Information Element (IE)	ID	Size (Bytes)	Notes
reportId	12276 - 55	4	
observationTimeSeconds	12276 - 90	8	
timestampMsec	12276 - 91	2	
recordType	12276 - 54	1	
subscriberId	12276 - 71	Variable	This IE is omitted for NetFlow v9.
subscriberIdType	12276 - 72	Variable	This IE is omitted for NetFlow v9.
3gppParameters	12276 - 57	Variable	This IE is omitted for NetFlow v9.
applicationCategoryId	12276 - 48	2	
lastRecordSent	12276 - 63	8	
uplinkVolume	12276 - 89	8	
downlinkVolume	12276 - 88	8	
concurrentFlows	12276 - 59	2	
newFlows	12276 - 64	2	
terminatedFlows	12276 - 69	2	
totalTransactions	12276 - 73	2	
successfulTransactions	12276 - 68	4	
durationSec	12276 - 60	2	
recordReason	12276 - 66	1	
reportVersion	12276 - 56	Variable	This IE is omitted for NetFlow v9.

Flow logs

This IPFIX template is used for flow records used for HSL reporting.

Information Element (IE)	ID	Size (Bytes)	Notes
reportId	12276 - 55	4	
observationTimeSeconds	12276 - 90	8	
timestampMsec	12276 - 91	2	
recordType	12276 - 54	1	
subscriberId	12276 - 71	Variable	This IE is omitted for NetFlow v9.
subscriberIdType	12276 - 72	Variable	This IE is omitted for NetFlow v9.
sourceIPv4Address	8	4	

Information Element (IE)	ID	Size (Bytes)	Notes
sourceIPv6Address	27	16	
sourceTransportPort	7	2	
destinationIPv4Address	12	4	
destinationIPv6Address	28	16	
destinationTransportPort	11	2	
protocolIdentifier	4	1	
applicationCategoryId	12276 - 48	2	
urlCategoryId	12276 - 87	2	
flowStartSeconds	12276 - 51	8	
flowStartMilliSeconds	12276 - 50	2	
flowStopSeconds	12276 - 53	8	
flowStopMilliSeconds	12276 - 52	2	
totalTransactions	12276 - 73	2	
uplinkVolume	12276 - 89	8	
downlinkVolume	12276 - 88	8	
reportVersion	12276 - 56	Variable	This IE is omitted for NetFlow v9.
ingressVRFD	234	4	
vlanId	12276 - 92	2	

Transaction logs

This IPFIX template is used for transactional records used for HSL reporting.

Information Element (IE)	ID	Size (Bytes)	Notes
reportId	12276 - 55	4	
recordType	12276 - 54	1	
reportVersion	12276 - 56	Variable	This IE is omitted for NetFlow v9.
transactionNumber	12276 - 81	2	
subscriberId	12276 - 71	Variable	This IE is omitted for NetFlow v9.
subscriberIdType	12276 - 72	Variable	This IE is omitted for NetFlow v9.
sourceIPv4Address	8	4	
sourceIPv6Address	27	16	
sourceTransportPort	7	2	
destinationIPv4Address	12	4	
destinationIPv6Address	28	16	
destinationTransportPort	11	2	

Information Element (IE)	ID	Size (Bytes)	Notes
protocolIdentifier	4	1	
ingressVRFID	234	4	
vlanId	12276 - 92	2	
applicationCategoryId	12276 - 48	2	
urlCategoryId	12276 - 87	2	
classification	12276 - 49	Variable	This IE is omitted for NetFlow v9.
transactionStartSeconds	12276 - 84	8	
transactionStartMilliseconds	12276 - 83	2	
transactionStopSeconds	12276 - 86	8	
transactionStopMilliseconds	12276 - 85	2	
uplinkVolume	12276 - 89	8	
downlinkVolume	12276 - 88	8	
skippedTransactions	12276 - 82	2	
httpResponseCode	12276 - 76	2	
httpHostnameTruncated	12276 - 75	1	
httpHostname	12276 - 74	Variable	This IE is omitted for NetFlow v9.
httpUserAgentTruncated	12276 - 80	1	
httpUserAgent	12276 - 79	Variable	This IE is omitted for NetFlow v9.
httpUrlTruncated	12276 - 78	1	
httpUrl	12276 - 77	Variable	This IE is omitted for NetFlow v9.

Legal Notices

Legal notices

Publication Date

This document was published on May 9, 2016.

Publication Number

MAN-0404-07

Copyright

Copyright © 2012-2016, F5 Networks, Inc. All rights reserved.

F5 Networks, Inc. (F5) believes the information it furnishes to be accurate and reliable. However, F5 assumes no responsibility for the use of this information, nor any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent, copyright, or other intellectual property right of F5 except as specifically described by applicable user licenses. F5 reserves the right to change specifications at any time without notice.

Trademarks

For a current list of F5 trademarks and service marks, see
<http://www.f5.com/about/guidelines-policies/trademarks/>.

All other product and company names herein may be trademarks of their respective owners.

Patents

This product may be protected by one or more patents indicated at:
<http://www.f5.com/about/guidelines-policies/patents>

Export Regulation Notice

This product may include cryptographic software. Under the Export Administration Act, the United States government may consider it a criminal offense to export this product from the United States.

RF Interference Warning

This is a Class A product. In a domestic environment this product may cause radio interference, in which case the user may be required to take adequate measures.

FCC Compliance

This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This unit generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area

is likely to cause harmful interference, in which case the user, at his own expense, will be required to take whatever measures may be required to correct the interference.

Any modifications to this device, unless expressly approved by the manufacturer, can void the user's authority to operate this equipment under part 15 of the FCC rules.

Canadian Regulatory Compliance

This Class A digital apparatus complies with Canadian ICES-003.

Standards Compliance

This product conforms to the IEC, European Union, ANSI/UL and Canadian CSA standards applicable to Information Technology products at the time of manufacture.

Index

A

- active sessions
 - records [151](#)
 - statistics [151](#)
- all-match strategy
 - about [133](#)
- Any IP profile
 - updating [144](#)
- Any IP profiles
 - creating [143](#)
- application rules
 - adding to enforcement policy [23](#), [38](#), [108](#)
- applications
 - creating custom classification [16](#), [128](#)
 - overview custom classification [127](#)
- application visibility
 - defined [15](#)
- application visibility statistics
 - examining [18](#)

B

- bandwidth control
 - creating a rule for [81](#)
- bandwidth controller categories
 - adding [89](#)
- bandwidth controllers
 - compared with rate shaping [85](#)
- bandwidth control policies
 - about Max User Rate setting [87](#)
 - adding categories to [89](#)
 - and SNMP [85](#)
 - creating dynamic for PCRF [91](#)
 - dynamic, about [87](#)
 - dynamic, adding to virtual server [90](#)
 - dynamic, classifying traffic [89](#)
 - dynamic, creating [88](#)
 - dynamic, prerequisites [87](#)
 - for tiered services [95](#)
 - overview [85](#), [95](#)
 - static, about [85](#)
 - static, adding to virtual server [86](#)
 - static, creating [80](#), [86](#)
- bandwidth control with PEM
 - overview [79](#), [91](#)
 - result of [93](#)
- bandwidth control with tiered services
 - result of [99](#)
- best-match strategy
 - about [133](#)
- best practices, PEM [11](#)

C

- CARP
 - overview [19](#)

- categories
 - creating custom classification [16](#), [127](#)
 - determining classification [15](#), [127](#)
- classification
 - using iRules [129](#)
- classification applications
 - creating custom [16](#), [128](#)
 - overview [127](#)
- classification categories
 - creating custom [16](#), [127](#)
 - determining [15](#), [127](#)
- classification data
 - overview [15](#)
- classification iRule commands [131](#)
- classification signatures
 - updating automatically [125](#)
 - updating overview [57](#), [125](#)
- classification statistics
 - examining [18](#)
- collectors
 - for IPFIX [41](#)
- content insertion
 - creating policy action [33](#)
- creating [103](#)
- custom action
 - iRules [22](#)

D

- destinations
 - for IPFIX logging [42](#)
- device detection
 - enabling tethering [49](#), [145–147](#)
 - overview [145](#)
- DHCPv4 profiles
 - creating [69](#)
- DHCPv6 profile
 - creating [71](#)
- Diameter
 - about [36](#)
 - configuring AVP [124](#)
- dynamic bandwidth control policies, See bandwidth control policies. [87](#), [89–90](#)
- dynamic service chains
 - creating [105](#)
- dynamic subscribers
 - overview [113](#)
 - provisioning [113](#)

E

- endpoints
 - creating [20](#)
 - creating for service chains [105](#)
- enforcement policy
 - about rules [10](#)
 - adding Gx reporting rules [76](#)
 - adding HSL reporting rules [46–48](#)

- enforcement policy (*continued*)
 - adding QoE reporting rules 60
 - adding rules to 23, 38, 108
 - best practices 11
 - classifying traffic rules 24, 26
 - controlling bandwidth 81, 97
 - creating 22, 81, 96, 106
 - creating listeners 17, 30, 82, 110
 - overview 9–10
 - QoS traffic rules 28
 - troubleshooting 149
 - URL categorization 25
- enforcement rules
 - overview 10

F

- fastl4
 - creating 57–58
- first-match strategy
 - about 133
- forwarding traffic
 - creating rules 27, 109
 - overview 19

G

- global application policies
 - overview 79
- gx
 - policing 119
- Gx
 - configuring messages 122
 - configuring profile 122
 - messages 124
- Gx reporting
 - overview 67, 75
- Gy support
 - about 35

H

- header modifications
 - overview 33
- header values
 - for HTTP requests 102
- HSL reporting
 - overview 45, 59
- HSL reporting format
 - flow-based 51
 - HSL DTOS-based 53
 - qoe-based 61
 - session-based 49
 - tethering-based 54
 - transaction-based 52
- HTTP request-header values 102
- HTTP virtual server
 - send HTTP requests 103

I

- ICAP adaptation 102
- ICAP content adaptation 102
- ICAP profiles
 - 102
 - assigning 104
- inserting content
 - overview 33
- internal virtual server
 - for response modification 103
- internal virtual servers
 - creating 104
- internal virtual server type
 - defined 102
- IPFIX
 - AFM template overview 153
 - and server pools 41
 - template for sending IPFIX data set records for session reporting 156
 - template for sending IPFIX data set records for transactional reporting 156–157
- IPFIX collectors
 - and destinations for log messages 42
 - and publishers for log messages 43
- IPFIX logging
 - and PEM 41
 - creating a destination 42
- IPFIX logging, overview 41
- IPOther filter
 - PEM actions 144
- IPsec traffic
 - classification 143
- iRule commands, classification 131
- iRules
 - and bandwidth control policies 89
 - using with traffic classification 129

L

- limitation, PEM 12
- listener
 - creating 39, 75, 92, 98
- listener for DHCPv4 discovery virtual
 - creating 68
- listener for DHCPv6
 - creating 70
- listener for RADIUS
 - creating 73
- listeners
 - connecting to a PCRF 17, 30, 110
 - creating 17, 30, 110
- local traffic policy
 - associating with virtual servers 137
 - classification 133–134
 - creating strategy list 136
 - modifying PEM rules 135, 137
 - SSL 136
- local traffic policy matching
 - about all-match strategy 133
 - about best-match strategy 133
 - about first-match strategy 133

- local traffic policy matching (*continued*)
 - about matching strategies 133
- logging
 - and destinations 42
 - and pools 41
 - and publishers 43, 45, 59

- logging format
 - HSL DTOS session-based 53
 - HSL flow-based 51
 - HSL qoe-based 61
 - HSL session-based 49
 - HSL tethering-based 54
 - HSL transaction-based 52

M

- maximum rate of throughput, See bandwidth control policies
- modify header
 - creating 107

N

- New virtual group
 - connecting to a PCRF 82
 - creating 82

O

- online charging
 - about 36
- online charging system (OCS)
 - connecting to 39

P

- password 63
- PCRF
 - enforcing bandwidth control 91
 - provisioning subscribers 11
- PEM
 - custom TAC DB 141
 - IANA IPFIX IEs for 153
 - IPFIX template for PEM flow logs 156
 - IPFIX template for PEM session logs 156
 - IPFIX template for PEM transaction logs 157
- performance, PEM 12
- policy
 - creating enforcement 22, 81, 96, 106
 - provisioning dynamic subscribers 113
- Policy and Charging Rules Function (PCRF)
 - connecting to 75, 92, 98
- policy enforcement
 - overview 9
- policy re-evaluation
 - interval 139
- policy updates
 - overview 139
- pools
 - creating 20, 104
 - for IPFIX 41

- profile
 - creating listeners 57–58
- publishers
 - and logging 43
 - creating for logging 45, 59

Q

- quota management
 - global 37
- quota management with PEM
 - overview 35

R

- radius
 - provision 119
- RADIUS
 - configuring profile 119
 - Profile 121
- RADIUS AAA profiles 63
- RADIUS AAA virtual
 - creating 64
- RADIUS accounting
 - report 64
- RADIUS authentication
 - overview 63
- RADIUS AVP
 - creating 120
- RADIUS message
 - ingress 121
- RADIUS re-transmission
 - overview 139
 - re-transmit timeout 139
- rate shaping
 - compared with bandwidth controllers 85
- rating group
 - configuring 37
- rating groups
 - about 35
 - creating 37
 - definition 36
- recommendations, PEM 11
- reporting format
 - HSL DTOS session-based 53
 - HSL flow-based 51
 - HSL qoe-based 61
 - HSL session-based 49
 - HSL tethering-based 54
- request adapt profile 103
- response adapt profile 103
- rule
 - creating 64
- rules
 - about enforcement 10
 - for bandwidth control 81, 97
 - for classifying traffic 24, 26
 - for QoS 28
 - local traffic policy 135, 137
 - setting up Gx reporting 76
 - setting up HSL reporting 46–48
 - setting up QoE reporting 60

S

- secret 63
- servers
 - and destinations for log messages 42
 - and publishers for IPFIX logs 43
 - and publishers for log messages 45, 59
- service chain
 - creating endpoints for 105
- service chain endpoint
 - process 102
- service chains
 - overview 101
- signatures
 - configuring updates 125
- SNMP
 - and bandwidth control policies 85
- static bandwidth control policies, *See* bandwidth control policies
- static subscribers
 - overview 115
- statistics
 - examining for classification 18
- steering
 - adding classification rules to policy 23, 38, 108
 - creating a policy 22, 81, 106
 - creating endpoints 20
 - creating service chains 105
- steering policy
 - adding 107
 - define 102
- subscriber activity
 - debuggability 152
- subscriber attributes
 - configuring profile 123
- subscriber CSV file format 116
- subscriber identity 119
- subscribers
 - adding 115
 - overview 113, 115
 - provisioning dynamic 113
 - provisioning with PCRF 11
 - uploading 117

T

- tac db
 - custom configuration 141
- TAC DB
 - overview custom creating 141
- tethering
 - sending reports 49, 145–147
- tethering, PEM
 - result of 148
- tiered services
 - overview 95
- traffic classification
 - defined 15
 - overview 15
 - using iRules 129
- Traffic Intelligence
 - applications 127
- traffic steering
 - overview 19
- transactional policy enforcement
 - overview 19
- troubleshooting, PEM 149

U

- URL categories
 - creating iRule Events 26, 130
- URL database
 - creating custom 128
- URL Filtering
 - creating custom classification 128

V

- video Quality of Experience
 - creating virtual server 60
- virtual servers
 - adding dynamic bandwidth control policy 90
 - adding static bandwidth control policies 86
 - and internal type 102
 - associating local traffic policy 137
 - creating 104
 - creating for video Quality of Experience 60
- VLANs
 - creating 79